Intel® NetStructure™ 480T Routing Switch

User Guide

Intel® NetStructure™ 480T Routing Switch

User Guide

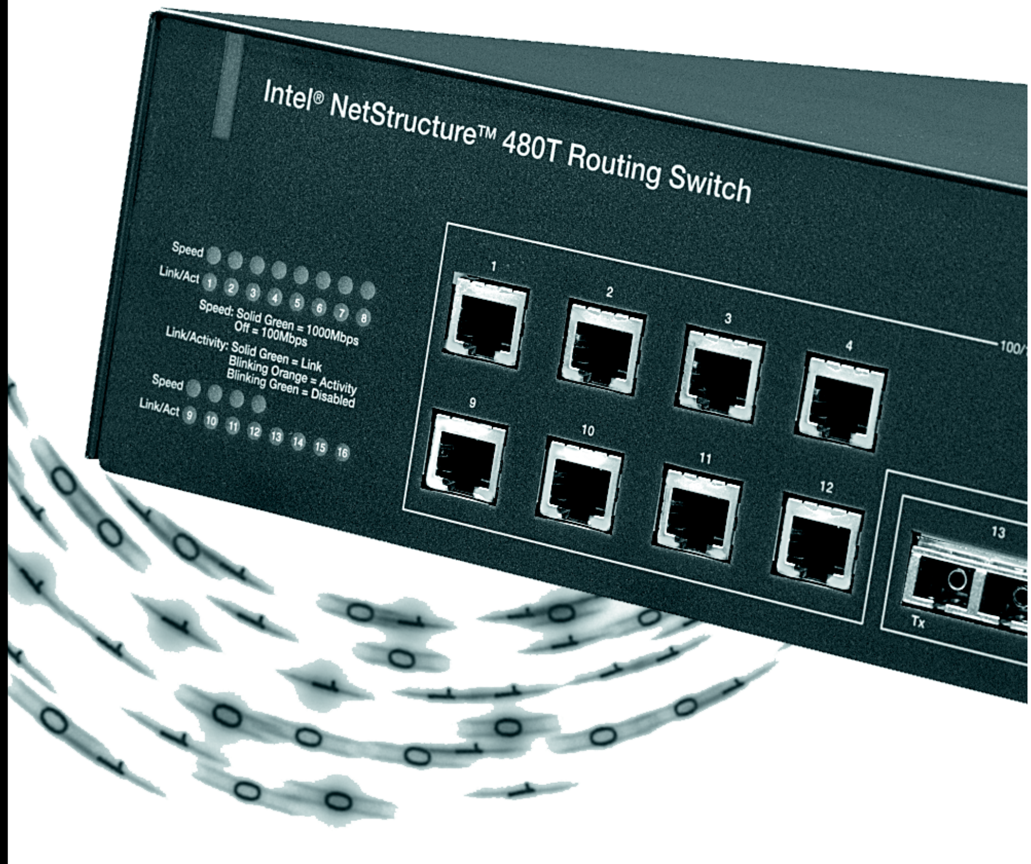

Intel® NetStructure™ 480T Routing Switch
Speed
Link/Act
1 2 3 4 5 6 7 8
Speed: Solid Green = 1000Mbps
Off = 100Mbps
Link/Activity: Solid Green = Link
Blinking Orange = Activity
Blinking Green = Disabled
Speed
Link/Act
9 10 11 12 13 14 15 16
1
2
3
4
9
10
11
12
13
Tx

intel®

Intel Corporation, 5200 NE Elam Young Parkway, Hillsboro OR 97124-6497

Intel Corporation assumes no responsibility for errors or omissions in this manual. Nor does Intel make any commitment to update the information contained herein.

FourthEdition November 2001 A14542-001

Contents

2: Installation and Setup 17

3: Using Intel® Device View 23

##

##

19: Status Monitoring and Statistics 403

20: Software Upgrade and Boot Options 419

List of Figures

List of Tables

Preface

This preface provides an overview of this user guide, describes guide conventions, and lists other useful publications.

Introduction

This user guide provides the information you need to configure the Intel® NetStructure™ 480T routing switch.

Information in the "Late Breaking News" shipped with your switch is more up to date than the information in this guide.

It is intended for use by network administrators who are responsible for installing and setting up network equipment, and assumes a basic working knowledge of:

- Local Area Networks (LANs)
- Ethernet concepts, including switching and bridging
- Routing
- Internet Protocol (IP)
- Routing Information Protocol (RIP) and Open Shortest Path First (OSPF)
- Border Gateway Protocol (BGP-4)
- IP Multicast
- Distance Vector Multicast Routing Protocol (DVMRP)
- Protocol Independent Multicast (PIM)

- Internet Packet Exchange (IPX)
- Server Load Balancing (SLB)
- Simple Network Management Protocol (SNMP)

Related Publications

For further information refer to these publications:

- *Command Line Interface Reference Guide*
- *Intel® NetStructure™ 480T Routing Switch Quick Start Guide*
- *Late Breaking News*

Documentation for Intel products is available on the World Wide Web at the Intel support home page:

`http://support.intel.com`

1 Overview

The Intel® NetStructure™ 480T routing switch uses a powerful, full-featured software operating system for local management of the switch. This chapter offers an overview of the switch operation and covers these topics:

- Summary of features
- Software licensing
- Hardware specifications and factory defaults
- Media types

Summary of Features

The features of the 480T routing switch include:

- Virtual local area networks (VLANs) including support for IEEE 802.1Q and IEEE 802.1p (priority queuing)
- VLAN aggregation
- Spanning Tree Protocol (STP) (IEEE 802.1D) with multiple STP domains
- Policy-Based Quality of Service (PB-QoS)
- Wire-speed IP routing

- IP Multinetting
- Dynamic Host Configuration Protocol (DHCP)/Bootstrap Protocol (BOOTP) Relay
- Enterprise Standby Router Protocol (ESRP)
- RIP (Routing Information Protocol) version 1 and version 2
- OSPF (Open Shortest Path First) routing protocol
- BGP-4
- Wire-speed IP multicast routing support
- Diffserv (Differentiated Services) protocol support
- Access policy support for routing protocols
- Access list support for packet filtering
- IGMP (Internet Group Management Protocol) snooping to control IP multicast traffic
- DVMRP (Distance Vector Multicast Routing Protocol)
- Protocol Independent Multicast-Dense Mode (PIM-DM)
- Protocol Independent Multicast-Sparse Mode (PIM-SM)
- Wire-speed IPX[§], IPX/RIP, and IPX/Service Advertising Protocol (SAP) support
- SLB support
- Load sharing (link aggregation) on multiple ports
- RADIUS (Remote Authorization Dial-In User Service) client and per-command authentication support
- TACACS+ (Terminal Access Controller Access Control System) support
- Console command line interface (CLI) connection
- Telnet CLI connection
- Web-based management interface
- Simple Network Management Protocol (SNMP) support
- RMON (Remote Monitoring)
- Traffic mirroring for all ports
- Intel® Device View (IDV) support

Full-Duplex Support

The 480T routing switch provides full-duplex support for all ports. Full-duplex mode allows frames to be transmitted and received simultaneously and, in effect, doubles the bandwidth available on a link. All 100/1000 Mbps ports on the 480Tswitch autonegotiate for half-duplex or full-duplex operation.

The 1000BASE-SX, 1000BASE-LX and 1000LH ports operate in full-duplex mode only.

Virtual LANs (VLANs)

The local management software has a VLAN feature that enables you to construct your broadcast domains without being restricted by physical connections. A VLAN is a group of location and topology-independent devices that communicate as if they were on the same physical LAN.

Implementing VLANs on your network has three advantages:

- Better broadcast traffic control - If a device in VLAN ***Marketing*** transmits a broadcast frame, only VLAN ***Marketing*** devices receive the frame.

See Chapter 7, "Virtual LANs (VLANs)" on page 95.

- Extra security - Devices in VLAN ***Marketing*** can only communicate with devices in VLAN ***Sales*** using routing services.
- Easier to change or move devices on your networks.

Spanning Tree Protocol (STP)

The 480T routing switch supports the IEEE 802.1D Spanning Tree Protocol (STP), a bridge-based method of providing fault tolerance on networks. STP enables you to implement parallel paths for network traffic, and ensure that redundant paths are:

See Chapter 9, "Spanning Tree Protocol (STP)" on page 125.

- Disabled when the main paths are operational.
- Enabled if the main traffic paths fail.

A single spanning tree may span multiple VLANs.

Quality of Service (QoS)

See Chapter 10,"Quality of Service (QoS)" on page 135.

The local management software has Policy-Based Quality of Service (QoS) features that enable you to specify service levels for different traffic groups. By default, all traffic is assigned a normal QoS policy profile.

You can create other QoS policies and apply them to different traffic types so that they have different guaranteed minimum bandwidth, maximum bandwidth, and priority.

Unicast Routing

The 480T routing switch can route IP or IPX traffic between VLANs that are configured as virtual router interfaces. Both dynamic and static IP routes are maintained in the routing table. The routing protocols supported include:

See “IP Unicast Routing” on page 189.

- RIP version 1
- RIP version 2
- OSPF-2
- IPX/RIP
- BGP-4

For further information consult these chapters:

- "IP Unicast Routing" on page 189
- "RIP and OSPF" on page 223
- "Border Gateway Protocol (BGP)" on page 255
- "IPX Routing" on page 291

IP Multicast Routing

See “IP Multicast Routing” on page 275.

The 480T routing switch enables you to use IP multicasting to allow a single IP host to transmit a packet to a group of IP hosts. It supports multicast routes learned by way of the Distance Vector Multicast Routing Protocol (DVMRP) or Protocol Independent Multicast, dense or sparse mode (PIM-DM or PIM-SM).

Load Sharing

See "Configuring Ports" on page 79.

Load sharing allows you to increase bandwidth and resiliency by using a group of ports to carry traffic in parallel between systems. The switch's sharing algorithm allows you to use multiple ports as a single logical port.

For example, VLANs treat the load-sharing group as a single virtual port.

Software Licensing - Router License Keys

You can expand the feature set of your switch using a license key. The keys are unique to the 480T routing switch and are not transferable. Keys are stored in NVRAM and, once entered, persist through reboots, software upgrades, and later reconfigurations.

In the firmware, routing protocol support is separated into two sets:

- Basic
- Full Layer 3.

Basic is a subset of Full Layer 3.

Basic Functionality

Basic functionality requires no license key. It includes all switching functions, as well as all available Layer 3 QoS, access list, and ESRP functions.

Basic includes support for these Layer 3 routing functions:

- IP routing using RIP version 1, RIP version 2, or both
- IP routing between directly attached VLANs
- IP routing using static routes

Full Layer 3 Functionality

Switches using a Full Layer 3 license also support other routing protocols and functions in addition to Basic functions, including:

- IP routing using OSPF
- IP multicast routing using DVMRP
- IP multicast routing using PIM (Dense or Sparse Mode)
- IPX routing (direct, static, and dynamic using IPX/RIP and IPX/SAP)
- IP routing using BGP
- Server load balancing (SLB)
- Web cache redirection

Verifying the Router License

To verify the router license, use the **`show switch`** command.

Upgrading a Router License

You can upgrade the router license of a switch by purchasing a voucher from Intel. The voucher contains instructions on obtaining a license key from the Intel web site at support.intel.com.

Once a license key is entered, it is not necessary to enter the information again. We recommend keeping the upgrade voucher for your records.

Physical Features

Front View

Figure 1.1 shows the switch front view.

The 480T routing switch has 12 100/1000-Mbps ports, and four 1000 Mbps-only ports. Ports 13 through 16 use modular GBIC connectors.

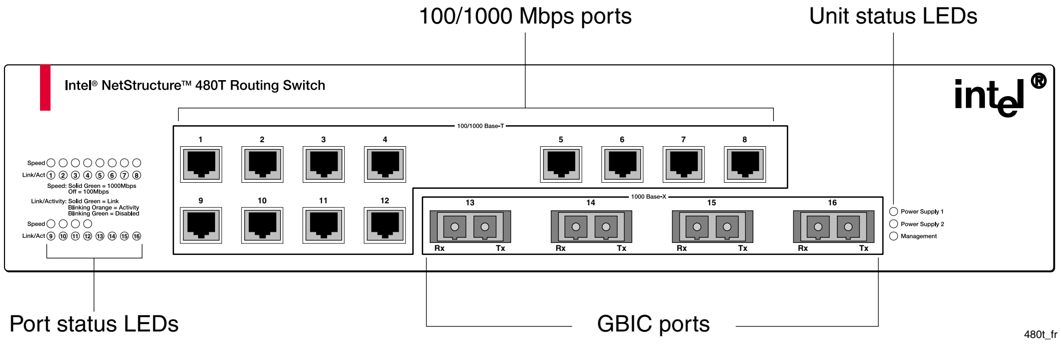

Figure 1.1: Intel® NetStructure™ 480T routing switch (front)

For information on switch LEDs, refer to "Switch LEDs" on page 10.

Rear View

Figure 1.2 shows two rear view configurations. The second has a redundant power supply.

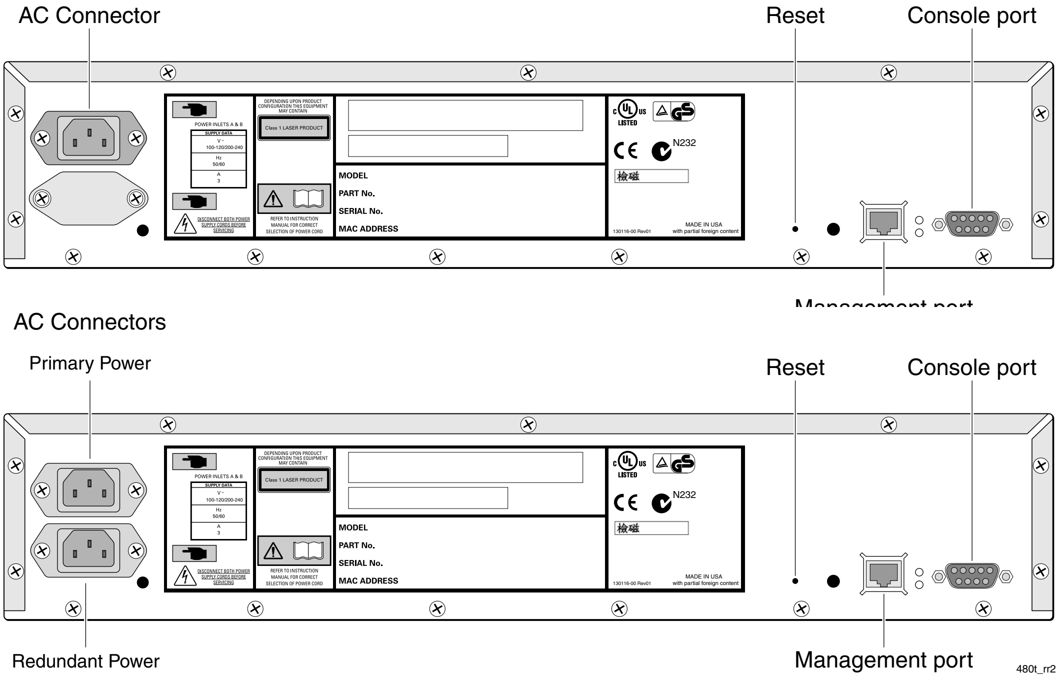

Figure 1.2: Intel® NetStructure™ 480T routing switch (with and without redundant power supply)

AC Connector

The 480T routing switch automatically adjusts to the supply voltage. The power supply unit (PSU) operates down to 100V, and is suitable for both 110 VAC and 200-240 VAC operation.

Serial Number

Use this serial number for fault-reporting purposes.

Console Port

Use the console port (9-pin, D-type connector) for connecting a terminal and carrying out local out-of-band management.

For information on supported media types and distances, refer to Table 1.3 on page 14.

Management Port

The management port (RJ-45 connector) is a 10/100 Mbps Ethernet connection used for out-of-band management.

MAC Address

This label shows the unique Ethernet MAC address assigned to this device.

Switch LEDs

Table 1.1 describes the light emitting diode (LED) behavior on the 480T routing switch.

Table 1.1: Switch LEDs

LED	Color	Indicates
1000BASE-X Port Status LEDs (GBIC LEDs)		
Link/activity	Green	Link is present; port is enabled.
	Orange	Frames are being transmitted/received on this port.
	Green flashing (steady)	Link is present; port is disabled.
	Off	Link is not present.
100/1000BASE-T Port Status LEDs		
Link/activity	Green	Link is present; port is enabled.
	Orange	Frames are being transmitted/received on the port.
		Link is present; port is disabled.
	Green flashing (steady)	Link is not present.
	Off	
Speed Status	Green	1000 BASE-T operation.
	Off	100 BASE-TX operation.
10/100 Management Port Status LEDs		
Link/activity	Green	Link is present.
	Orange	Frames are passing through this port.
	Off	Link is not present.
Unit Status LEDs		
Power 1 and Power 2	Green	Either or both LEDs green indicates the 480T routing switch is powered up.
	Orange	An orange power LED indicates a power, overheat, or fan failure on the corresponding power supply unit.
	Off	Both LEDs off indicates the switch is powered off.
MGMT	Green flashing (slow)	The 480T routing switch is operating normally.
	Green flashing (fast)	POST is in progress.
	Orange	The switch has failed POST.

Software Factory Defaults

Table 1.2 lists factory defaults for global features.

Table 1.2: Global Factory Defaults

Item	**Default Setting**
Serial or Telnet user account	`admin` with no password and `user` with no password
Web network management	Enabled
Telnet	Enabled
SNMP access	Enabled
SNMP read community string	`public`
SNMP write community string	`private`
RMON	Enabled
BOOTP	Enabled on the default VLAN
Quality of Service (QoS)	Disabled. If enabled, all traffic is part of the default queue
QoS monitoring	Automatic roving
802.1p priority	Recognition enabled
802.3x flow control	Enabled on 1000 Mbps Ethernet ports
CLI idle timeout	Enabled (15 minutes)
Virtual LANs	Three VLANs pre-defined. VLAN named default contains all ports and belongs to the STPD named s0. VLAN **mgmt** operates on the 10/100 Ethernet management port. The management port is DTE only, and is not capable of switching or routing. VLAN ***MacVLanDiscover*** is active only when using MAC VLAN.

Table 1.2: Global Factory Defaults (continued)

Item	Default Setting
802.1Q tagging	Packets are untagged on the default VLAN.
Spanning Tree Protocol	Disabled for the Intel® NetStructure™ 480T routing switch; enabled for each port in the STPD
Forwarding database aging period	300 seconds (5 minutes)
IP Routing	Disabled
RIP	Disabled
OSPF	Disabled
IP multicast routing	Disabled
IGMP	Enabled
IGMP snooping	Enabled
DVMRP	Disabled
PIM	Disabled
IPX§ routing	Disabled
NTP	Disabled
DNS	Disabled
Port mirroring	Disabled
Server load balancing	Disabled
Web Cache Redirection	Disabled
ESRP	Disabled
BGP-4	Disabled

Media Types, Distances and Specifications

Table 1.3 describes the media types and distances (cable lengths) for the different types of switch ports.

Table 1.3: Media Types and Distances

Type	Media	M Hz/Km Rating	Maximum Distance
1000BASE-SX	50/125 µm Multimode Fiber	400	500 Meters
	50/125 µm Multimode Fiber	500	550 Meters
	62.5/125 µm Multimode Fiber	160	220 Meters
	62.5/125 µm Multimode Fiber	200	275 Meters
1000BASE-LX	50/125 µm Multimode Fiber	400	550 Meters
	50/125 µm Multimode Fiber	500	550 Meters
	62.5/125 µm Multimode Fiber	500	550 Meters
	10µ Single-mode Fiber		5 Kilometers
1000LH	10µ Single-mode Fiber		70 Kilometers
1000BASE-T	Category 5 and higher UTP Cable		100 Meters
100BASE-TX	Category 5 and higher UTP Cable		100 Meters
10BASE-T	Category 3 and higher UTP Cable		100 Meters

Table 1.4 describes the specifications for the 1000B-LH interface.

Table 1.4: 1000LH Specifications

Parameter	Minimum	Typical	Maximum
Transceiver			
Optical Output Power	0 dBm	3 dBm	5 dBm
Center Wavelength	1540 nm	1550 nm	1560 nm
Receiver			
Optical Input Power Sensitivity	-20 dBm		
Optical Input Power Maximum			-3d Bm
Operating Wavelength	1200nm		1560 nm

Optical Output Power

The minimum cable length without a 10 dB attenuator is 32 kilometers.

The transmitter output power level for the 1000-LH is +5dBm. The maximum allowable receiver input power level is -3dBm. Therefore, there is a minimum of 8dB loss required for the link to operate without errors. You can achieve this minimum required loss using a fiber length of 32km (0.25dB/km provides 8dB loss), or by adding 10dB of fixed optical attenuator at the receiver end.

2 Installation and Setup

This chapter describes:

- Determining the Switch Location
- Installing the Switch
- Connecting Equipment to the Console Port
- Checking the Installation Using the Power-On Self Test (POST)
- Logging In for the First Time
- Upgrading Your Firmware
- Installing the Gigabit Interface Connector (GBIC)

Important Safety Information

Safety related specifications are provided in Appendix A, "Technical Specifications and Supported Limits" on page 431.

There are no user serviceable parts on the Intel® NetStructure™ 480T routing switch. The switch uses Class 1 laser technology. The ports emit invisible infrared light. Do not look directly into open ports.

Determining the Switch Location

The 480T routing switch can be free standing or mounted in a standard 19-inch equipment rack. Mounting brackets are supplied with the switch.

When deciding where to install the switch, ensure that:

- The switch is accessible and you can connect cables easily.
- Water or moisture cannot enter the case of the unit.
- Air flow around the unit and through the side vents is not restricted.
- The switch has a minimum of 25 mm (1-inch) clearance.
- Units are not stacked more than four high if the switch is free-standing.

Installing the Switch

You can mount the switch in a rack or place it free-standing on a tabletop.

***Caution:** Do not suspend the switch from under a table or desk, or attach it to a wall.*

Rack Mounting

To rack mount the 480T routing switch:

1 Place the switch upright on a hard flat surface, with the front facing you.

2 Remove the screws (4 each side) from the sides of the chassis and retain for Step 4.

3 Place the mounting bracket over the mounting holes on one side of the unit.

4 Replace the screws and fully tighten with a screwdriver, as shown in Figure 2.1.

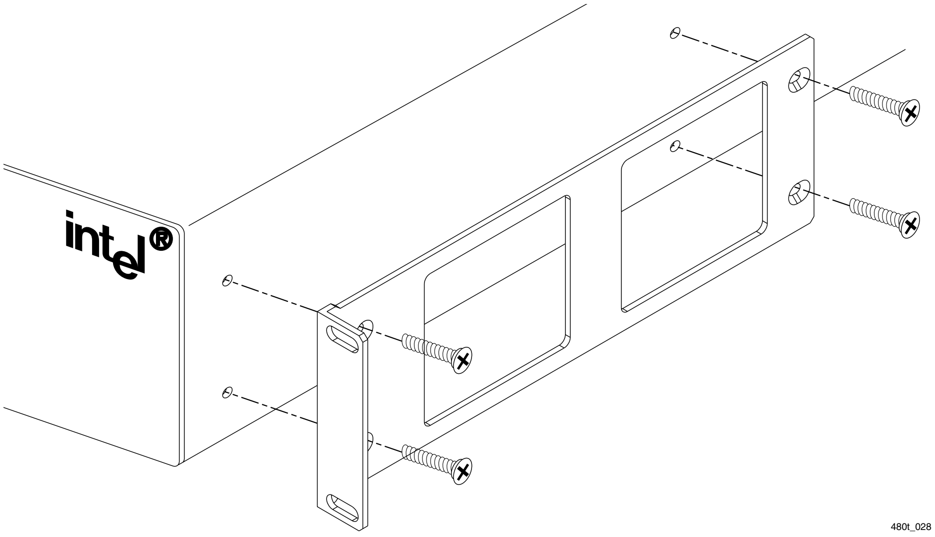

Figure 2.1: Fitting the mounting bracket

5 Repeat the two previous steps for the other side of the switch.

6 Insert the switch into the 19-inch rack. Ensure that ventilation holes are not obstructed.

7 Secure the switch with rack mount screws (not provided).

8 Remove the label over the AC connector and attach the power cord.

9 Attach the cables according to your own network configuration.

Many performance problems are caused by improper cabling. Pay careful attention to distance and cable type restrictions. See “Media Types, Distances and Specifications” on page 14.

Free-Standing

The 480T routing switch is supplied with four self-adhesive rubber pads.

You can stack up to four switches on top of one another.

1. Apply the pads to the underside of the device by sticking a pad in the marked area at each corner of the switch.
2. Place the devices on top of one another, ensuring that the corners align.

Connecting Equipment to the Console Port

For direct local management, connect to the console port. The 480T routing switch console port settings are set as follows:

- **Baud rate**—9600
- **Data bits—**8
- **Stop bit—**1
- **Parity**—None
- **Flow control**—XON/XOFF

Be sure the terminal connected to the console port on the switch is configured with the same settings. This procedure is described in the documentation supplied with the terminal or terminal emulation software.

Turning On the Switch

To turn on power to the switch, connect the AC power cable to the switch and then to the power outlet. The switch has no on/off switch.

Checking the Installation

After plugging in the switch, the device performs a Power-On Self-Test (POST).

During the POST, all ports are temporarily disabled, the packet LED is off and the power LED is on. The MGMT LED flashes quickly until the switch has successfully passed the POST, whereby it returns to the slow flashing state for normal operation.

If the switch passes the POST, the MGMT LED blinks at a slow rate (1 blink per second). If the switch fails the POST, the MGMT LED shows a solid orange light.

Logging In for the First Time

After the switch has completed the Power-On Self Test (POST), it is operational. Then you can log in to the switch and configure an IP address for the default VLAN (named ***default***).

To manually configure the IP settings:

1 Connect a terminal or workstation running terminal-emulation software to the console port.

2 At your terminal, press Enter one or more times until you see the login prompt.

3 At the login prompt, enter the default user name ***admin*** to log in with administrator privileges.

Administrator capabilities allow you to access all switch functions.

4 At the password prompt, press Enter.

The default name ***admin*** has no password assigned. When you have successfully logged in, the command-line prompt displays the name of the switch (for example, `Switch480T`) in its prompt.

5 Assign an IP address and subnetwork mask for VLAN ***default***. Use these commands (example IP addresses are used):

```
configure vlan default ipaddress 123.45.67.8
255.255.255.0
```

```
configure iproute add default <gateway>
123.45.67.8
```

```
enable ipforwarding
```

```
enable rip
```

Your changes should take effect immediately.

6 Save your configuration changes so that they are in effect after the next switch reboot. Use this command to save:

```
save
```

7 When you have finished, log out of the switch using this command:

```
logout
```

Upgrading Your Firmware

To upgrade your Intel® NetStructure™ 480T routing switch you must upgrade the BootRom image and firmware. Refer to the *Late Breaking News* that shipped with your switch for this procedure.

Installing the Gigabit Interface Connector (GBIC)

Ensure that the SC fiber-optic connector is removed from the GBIC prior to removing the GBIC from the I/O module.

You can add and remove Gigabit Interface Connectors (GBICs) from the 480T routing switch without powering off the system. Three types of GBIC modules are available:

- 1000BASE-SX
- 1000BASE-LX
- 1000LH

Warning: *Avoid exposing your eye to Class I laser radiation from open 1000 Mbps ports. Laser radiation is invisible to the human eye. Do not look directly into the 1000 Mbps port when installing or removing GBICs to eliminate any possible harmful effects. Class I lasers are not considered harmful under normal operation.*

Figure 2.2 illustrates a typical GBIC.

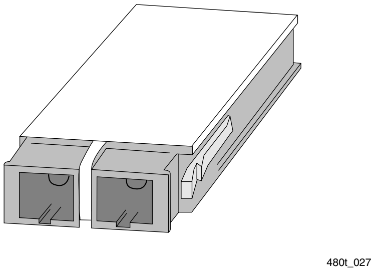

Figure 2.2: GBIC module (1000 Mbps ports)

GBICs are a Class 1 laser device. Use only Intel approved modules.

3 Using Intel® Device View

Intel® Device View is a graphical user interface that helps you manage the Intel NetStructure™ 480T routing switch and other supported Intel networking devices on your network.

Intel Device View provides these features:

- The ability to configure new network devices
- A graphical device manager for Intel switches, hubs, and routers
- Autodiscovery, which finds supported Intel devices on the network
- Device Tree, which shows all supported devices detected on your network
- Remote Network Monitoring (RMON)
- Web or Windows§ platform
- Plug-in to HP OpenView§, IBM Tivoli NetView§, and Intel LANDesk® Network Manager
- Other useful tools such as a TFTP server

Installing Intel Device View

Before you install Intel Device View, make sure your PC meets the system recommendations in the *Intel Device View User Guide*, which is included on the Intel Device View CD-ROM.

You can install both the Windows and the Web version of Intel Device View.

To Install Intel Device View

If you manage devices with Intel Device View from only one location on the network, install the Windows§ version.

1. Put the Intel Device View CD-ROM in your computer's CD-ROM drive. The Intel Device View installation screen appears. If it does not appear, run autoplay.exe from the CD-ROM (use the Run dialog from the Start menu).

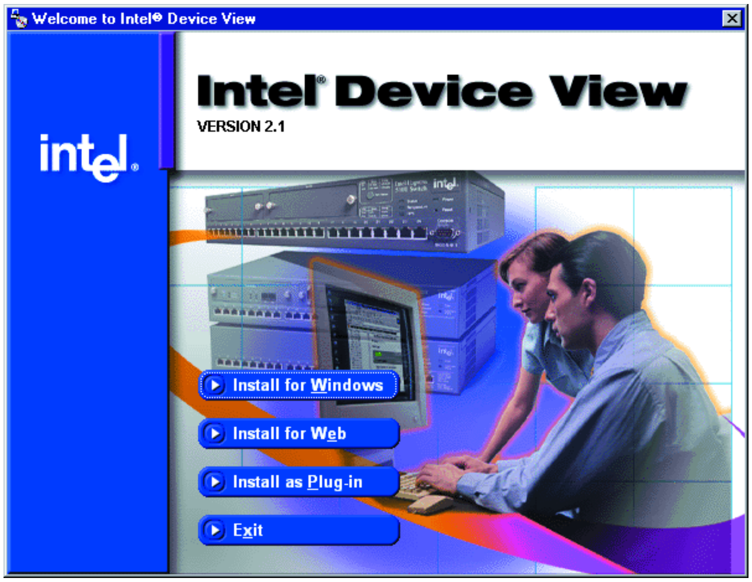

If you want to manage devices from any PC on the network using Intel Device View, install the Web version.

2. Choose the version of Intel Device View you want to install:
   - Click Install for Windows to install Intel Device View for use on this PC only.
   - Click Install for Web to install Intel Device View on a Web server. You is able to access the Device View server from any PC on your network with Internet Explorer§ 4.0x or later.
   - Click Install as Plug-in to install Intel network device support for HP OpenView, IBM Tivoli NetView, or Intel LANDesk Network Manager. This option is not available if you do not have any of these programs installed on the PC.
3. Follow the on-screen instructions in the installation program.

Starting the Windows§ Version

We recommend you use the Window version of Intel Device View if you manage devices from only one location on the network.

To start the Windows version:

1 From your desktop, click Start.

2 Point to Programs > Intel Device View > Intel Device View - Windows.

Intel Device View's main screen appears.

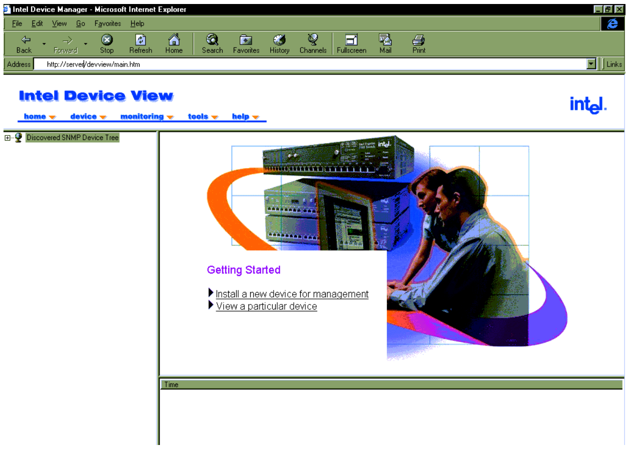

Address http://server/devview/main.htm

Intel Device Manager - Microsoft Internet Explorer
File Edit View Go Favorites Help
Back Forward Stop Refresh Home Search Favorites History Channels Fullscreen Mail Print
Links
Intel Device View
home device monitoring tools help
intel.
Discovered SNMP Device Tree
Getting Started
Install a new device for management
View a particular device
Time

Starting the Web Version

We recommend you use the Web version of Intel Device View if you want to manage devices from any PC on the network. To start the Web version:

1. From your desktop, click Start.
2. Point to Programs > Intel Device View > Intel Device View - Web. Intel Device View's main screen appears.

To view Intel Device View from another PC on your network, enter this URL into the Address field for Internet Explorer:

```
http://<servername>/devview/main.htm
```

where `<servername>` is the IP address or name of the server where Intel Device View is installed. Intel Device View's main screen appears.

Installing a New Device

After you've installed a new switch on your network, you can use Intel Device View's Device Install Wizard to configure it for management.

To Install and Configure a New Switch for Management

1. Start Intel Device View.

   The Device Install Wizard appears. If not, click Install from the Device menu or double-click the appropriate MAC address in the Device Tree under Unconfigured Devices.

2. In the Start screen, click Next.

3. In the MAC Address screen, click the MAC address of the new switch, and then click Next.

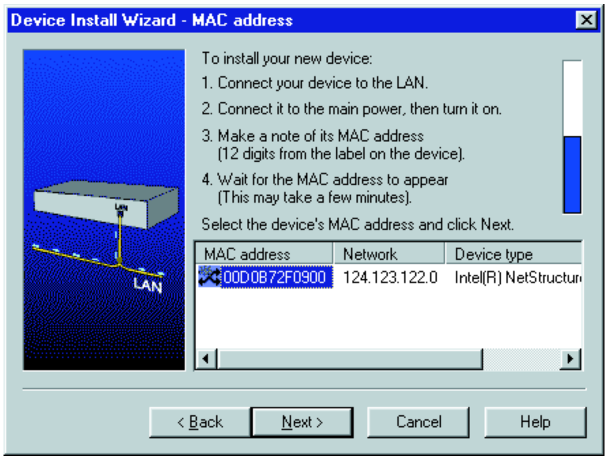

To install your new device:
1. Connect your device to the LAN.
2. Connect it to the main power, then turn it on.
3. Make a note of its MAC address
(12 digits from the label on the device).
4. Wait for the MAC address to appear
(This may take a few minutes).
Select the device's MAC address and click Next.

MAC address	Network	Device type
00D0B72F0900	124.123.122.0	Intel(R) NetStructur

4. Follow the instructions in the wizard to assign an IP address and a name to the switch.

Using the Device Tree

When you start Intel Device View, the Device Discovery service begins searching for supported Intel network devices on your network. As it discovers devices, it adds an icon for each device to the Device Tree on the left side of the screen.

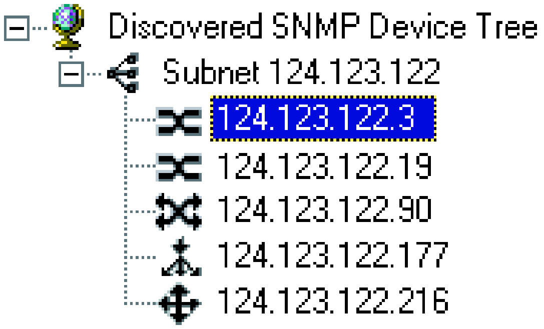

Discovered SNMP Device Tree
Subnet 124.123.122
124.123.122.3
124.123.122.19
124.123.122.90
124.123.122.177
124.123.122.216

Different states of the 480T routing switch are represented by unique icons in the Device Tree as indicated below.

Device Tree icons

Device Tree root

Subnet

Intel Switch (if non-responding the icon is red)

Unconfigured Intel Switch

Group of Intel Switches

Intel Router

Intel Switch (Layer 3 capable)

Intel Stackable Hub

The Device Tree works much like Windows Explorer:

- To expand the root or a subnet, click the (+) next to the icon.
- To collapse the view, click the (-) next to the icon.
- Double-click a device icon to view the device image.

To Add a Device to the Device Tree

1. Right-click anywhere on the Device Tree.
2. When a menu appears, click Add Device.
3. In the Add Device dialog box, enter the IP address of the switch you want to add.
4. Fill in the other fields, as appropriate.
5. Click OK.

The new switch's icon appears in the Device Tree.

To Refresh the Device Tree

1. Right-click anywhere on the Device Tree.
2. When a menu appears, click Refresh.

Refreshing the Device Tree updates it to show any newly discovered devices and changes in device status.

To Delete a Device from the Device Tree

1. Right-click the device you want to remove from the Device Tree.
2. Click Delete on the menu that appears.

Deleting a device from the Device Tree does not affect the actual device, but only removes the icon from the tree.

To Find a Device in the Device Tree

1. Right-click anywhere on the Device Tree.
2. When a menu appears, click Find.

3. In the Find Device dialog box, enter the IP address of the device you want to find in the tree.
4. Click OK.

The device’s icon is highlighted in the Device Tree.

Losing Contact with a Device

If Intel Device View loses contact with a switch, it replaces the switch icon with the red non-responding switch icon.

When the red non-responding switch icon appears, you will not be able to manage the device in Intel Device View.

If you’re unable to ping the device or start a Telnet session, try accessing the switch’s Local Management. See “Accessing the Switch” on page 39.

Managing a Switch

To manage a 480T routing switch, double-click the switch icon in the Device Tree. In the example shown below, the switch was assigned an IP address of 124.123.122.3.

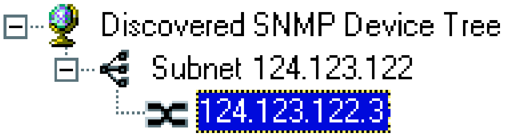

The Express 480T Web Device Manager appears in the Intel Device View window.

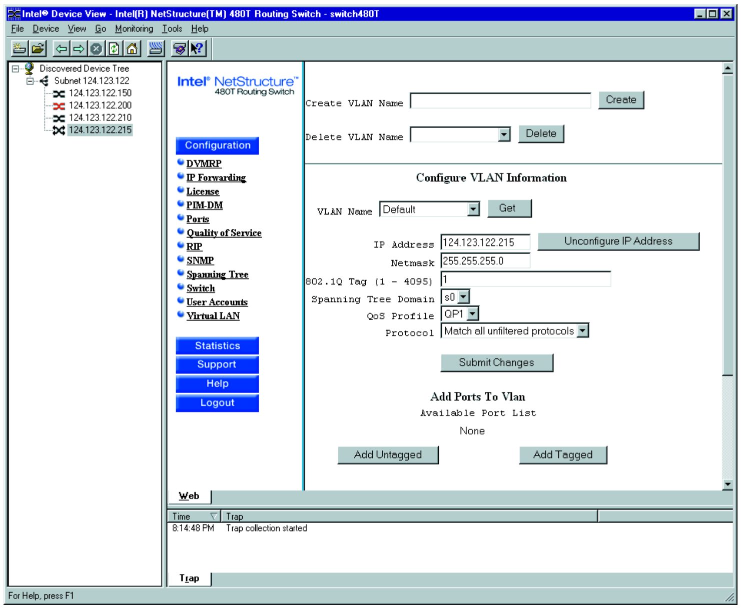

Create VLAN Name

Delete VLAN Name

Configure VLAN Information

VLAN Name: Default
IP Address: 124.123.122.215
Netmask: 255.255.255.0
802.1Q Tag (1 - 4095): 1
Spanning Tree Domain: s0
QoS Profile: QP1
Protocol: Match all unfiltered protocols

Add Ports To Vlan
Available Port List
None

Time	Trap
8:14:48 PM	Trap collection started

For complete information on using Intel Device View, refer to the program's online help or see the Intel Device View Help file on the installation CD-ROM.

Viewing RMON Information

The remote monitoring (RMON) specification is a feature of Intel Device View that extends Simple Network Management Protocol (SNMP) functionality to look at traffic patterns over the whole network instead of merely for an individual device. The 480T routing switch supports these RMON groups:

- **Group 1** Statistics**—**Monitors utilization and error statistics for each network segment (100Mbps or 1000Mbps).
- **Group 2** History**—**Records periodic statistical samples from variables available in the statistics group.
- **Group 3** Alarms**—**Allows you to set a sampling interval and alarm thresholds for statistics. When a threshold is passed, the

switch creates an event (see below). For example, you might set an alarm if switch utilization exceeds 30%.

- **Group 9** Events**—**Provides notification and tells the switch what to do when an event occurs on the network.

  Events can send a trap to a trap-receiving station, place an entry in the log table, or both. For example, when the switch experiences an RMON event, it sounds an alarm.

  The switch also keeps a log that shows a list of the RMON events and RMON alarms that have occurred on the switch.

To View RMON Statistics

1. In the Device Tree, right-click the switch's icon and then point to RMON.
2. Click the RMON option you want to view.

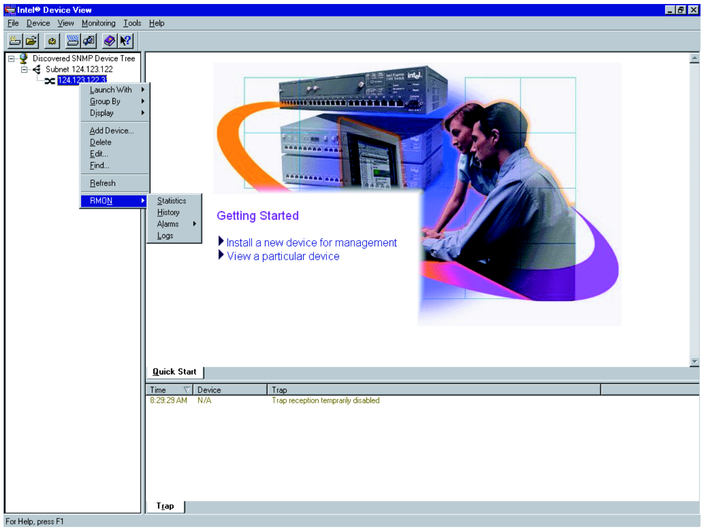

Time	Device	Trap
8:29:29 AM	N/A	Trap reception temprarily disabled

You can also access RMON features by using LANDesk Network Manager, or an SNMP application that supports RMON, such as OpenView.

For more information about using RMON to monitor the switch, refer to the Intel Device View Help file included on the CD-ROM.

4 Using Web Device Manager

Web Device Manager is device-management software running in the Intel® NetStructure™ 480T routing switch. It allows you to access the switch over a TCP/IP network, using a Web browser that supports frames and JavaScript§ (such as Netscape Navigator§ 3.0 or later, or Microsoft Internet Explorer§ 3.0 or later) to manage the system.

Web Device Manager provides a subset of the command-line interface (CLI) commands available for configuring and monitoring the switch. If a particular command is not available using Web Device Manager, use the CLI to access the desired functionality.

To use Web Device Manager, at least one VLAN must be assigned an IP address.

Enabling and Disabling Web Access

By default, Web access is enabled on the switch. You can restrict the use of Web access using an access profile.

For information on creating an access profile see page 324.

An access profile permits or denies a named list of IP addresses and subnet masks. To configure Web access to use an access profile, use this command:

```
enable web access-profile [<access-profile> | none]
{port <tcp_port_number>}
```

Use the `none` option to remove a configured access profile.

To display the status of Web access, use this command:

```
show management
```

To disable Web access, use this command:

```
disable web
```

To re-enable Web access, use this command:

```
enable web {access-profile [<access-profile> |
none]} {port <tcp_port_number>]
```

Reboot the system for these changes to take effect.

Setting Up Your Browser

Your browser's default settings should work well with Web Device Manager. Apply these recommended settings to improve the display features and functionality of Web Device Manager:

- After downloading a newer version of the switch image, clear the browser disk and memory cache to see the updated menu screens. It is important to clear the cache while at the main Logon screen, so that all underlying .GIF files are updated.
  - Check for newer versions of stored pages by setting the cache options to the "every visit" setting:
  - When using Netscape Navigator, configure the cache to check for changes Every Time you request a page.
  - When using Microsoft Internet Explorer, configure the Temporary Internet Files to check for newer versions of stored pages by selecting Every visit to the page.
- Images must be auto-loaded.
- Use a high-resolution monitor (1024 x 768 recommended) to maximize the amount of information displayed in the content frame. You can also use 800 x 600 pixels.
- Maximize viewing space by turning off the browser toolbars.
- Configure the browser to use these recommended fonts:
  - Proportional font—Times New Roman
  - Fixed-width font—Courier New

Accessing Web Device Manager

To access the default home page of the switch, enter this URL in your browser (substituting the actual ip address):

```
http://<ip_address>
```

When you access the home page of the system, the Login screen appears. Enter your user name and password and click OK.

If you have entered the name and password of an administrator-level account, you have access to all Web Device Manager pages. If you have used a user-level account name and password, you only have access to the Statistics and Support information.

If multiple people access the same switch using Web Device Manager, you might see this error message:

```
Web:server busy
```

To correct this situation, try logging out of the switch and logging in again.

Navigating Web Device Manager

After logging in to the switch, the Web Device Manager home page appears.

Web Device Manager divides the browser screen into these sections:

- Task frame
- Content frame
- Stand-alone buttons

Task Frame

The task frame has two sections: menu buttons and submenu links. There are four task menu buttons:

- Configuration
- Statistics
- Support
- Logout

Below the task buttons are options. Options are specific to the task button that you select. When you select an option, the information displayed in the content frame changes.

However, when you select a new task button, the content frame does not change until you select a new option.

Content Frame

When you submit a configuration page with no change an asterisk (*) will appear at the CLI prompt, even though actual configuration values have not changed.

The content frame contains the main body of information in Web Device Manager. For example, if you select an option from the Configuration task button, enter configuration parameters in the content frame. If you select the Statistics task button, statistics are displayed in the content frame.

Browser Controls

Browser controls include drop-down list boxes, check boxes, and multi-select list boxes. A multi-select list box has a scrollbar on the right side of the box. Using a multi-select list box, you can select a single item, all items, a set of contiguous items, or multiple non-contiguous items. Table 4.1 describes how to make selections from a multi-select list box.

Table 4.1: Multi-Select List Box Key Definitions

Selection Type	Key Sequence
Single item	Click the item using the mouse.
All items or contiguous items	Click the first item, and drag to the last item.
Contiguous items	Click the first item, hold down the Shift key, and click the last desired item.
Selected non-contiguous items	Hold down Ctrl, click the first desired item, click the next desired item, etc.

Status Messages

Status messages are displayed at the top of the content frame. There are four types of status messages:

- **Information**—Displays information that is useful to know prior to, or as a result of, changing configuration options.
- **Warning**—Displays warnings about the switch configuration.
- **Error**—Displays errors caused by incorrectly configured settings.
- **Success**—Displays informational messages after you click Submit. The message displayed reads, Request was submitted successfully.

Stand-alone Buttons

At the bottom of some of the content frames is a section that contains stand-alone buttons. Use these buttons to perform tasks that are not associated with a particular configuration option. An example of this is the Reboot Switch button.

Saving Changes

There are two ways to save your changes in Web Device Manager:

- Select Save Configuration from the Configuration task button, Switch option.

  This field contains a drop-down list box that allows you to select either the primary or secondary configuration area. After you select the configuration area, click Submit to save the changes.
- Click the Logout button.

  If you attempt to log out without saving your changes, Web Device Manager prompts you to save your changes.

  If you select Yes, the changes are saved to the selected configuration area.

To change the selected configuration area:

1. Go to the Configuration task button.
2. Select the Switch option.

Filtering Information

On some pages you can click a Filter button to display a subset of information for a page. For example, on the OSPF configuration page, you can configure authentication based on the VLAN, area identifier, or virtual link.

Once you select a filtering option and click the Filter button, the form that provides the configuration options displays the available interfaces in the drop-down menu, based on your filtering selection.

Using the *Get* Command to Configure a VLAN

When configuring a VLAN using Web Device Manager, prior to editing the VLAN configuration, you must first click the Get button to ensure that subsequent edits are applied to the correct VLAN. If you do not click the Get button and you submit the changes, the changes are made to the VLAN that was previously displayed.

If you configure a VLAN and then delete it, the ***default*** VLAN is shown in the VLAN name window, but the VLAN information contained in the lower portion of the page is not updated. Click the Get button to update the display.

TFTP Server

Intel Device View provides a TFTP Server utility on the Tools menu.

5 Accessing the Switch

This chapter provides information to help you manage the Intel® NetStructure™ 480T routing switch, including:

- Understanding the Command Syntax
- Line-Editing Keys
- Command History
- Common Commands
- Configuring Management Access
- Real-time Basic Connectivity Checking
- Methods of Managing the Switch
- Simple Network Management Protocol (SNMP))

*For information on using the **save** command, see "Software Upgrade and Boot Options" on page 419.*

To retain configuration changes through a power cycle or reboot, you must issue a `save` command after you have made the change.

Understanding the Command Syntax

This section briefly describes the steps to take when entering a command. The sections that follow give detailed information for using the command-line interface.

To use the command-line interface (CLI):

Most configuration commands require that you have administrator privileges.

1. Enter the command name.
   When entering a command at the prompt, ensure that you have the appropriate privilege level.
2. Enter the parameter name and values, if included.
   The value (also known as an argument) specifies how you want the parameter to be set. Values include numerics, strings, or addresses, depending on the parameter.
3. After entering the complete command, press Enter.

An asterisk (*) in front of the command-line prompt indicates you have made changes that have not been saved.

Syntax Helper

The CLI has a built-in syntax helper. If you are unsure of the complete syntax for a particular command, enter as much of the command as possible and press Enter. The syntax helper provides a list of options for the remainder of the command.

The syntax helper also provides assistance if you have entered an incorrect command.

Command Completion with Syntax Helper

Use the Tab key to access command completion.

1. Enter a partial command.
2. Press the Tab key to post a list of available options.
3. The cursor appears at the end of the command.

Abbreviated Syntax

Abbreviated syntax is the shortest, most unambiguous, allowable abbreviation of a command or parameter. Typically, this is the first three letters of the command. For example, *ena* is sufficient for the Enable command.

When using abbreviated syntax, you must enter enough characters to make the command unambiguous and distinguishable to the switch.

Command Shortcuts

All component names must be unique. Name components using the **`create`** command. When you enter a command to configure a named component, you do not need to use the keyword of the component. For example, to create a VLAN, you must enter a unique VLAN name:

```
create vlan engineering
```

After you create the VLAN with a unique name, you can eliminate the keyword **`vlan`** from all other commands that require the name to be entered. For example, instead of entering the command:

```
configure vlan engineering delete port 1,4
```

you can enter this shortcut:

```
configure engineering delete port 1,4
```

Numerical Ranges

Commands that require you to enter one or more port numbers on a switch use the parameter **`<portlist>`** in the syntax. For example:

```
port 3
```

A port list can be a range of numbers, for example:

```
port 1-3
```

You can add additional port numbers to the list, separated by a comma:

```
port 3,4,6
```

Names

All named components of the switch configuration must:

- Have a unique name.
- Begin with an alphabetical character.
- Be delimited (separated) by a space, unless enclosed in quotation marks.

Symbols

You may see a variety of symbols shown as part of the command syntax. These symbols explain how to enter the command, and you do not type them as part of the command itself. Table 5.1 summarizes command syntax symbols. Press the Tab key in the command line interface for more command options.

Table 5.1: Command Syntax Symbols

Symbol	Description
< > Angle brackets	Enclose a variable or value. You must specify the variable or value. For example, in the syntax: `configure vlan <name> ipaddress <ip_address>` you must supply a VLAN name for `<name>` and an address for `<ip_address>` when entering the command. Do not type the angle brackets.
[ ] Square brackets	Enclose a required value or list of required arguments. You can specify one or more values or arguments. For example, in the syntax: `use image [primary | secondary]` you must specify either the primary or secondary image when entering the command. Do not type the square brackets.
| Vertical bar	Separates mutually exclusive items in a list, one of which must be entered. For example, in the syntax: `configure snmp community [readonly | readwrite] <string>` you must specify either the read or the write community string in the command. Do not type the vertical bar.
{ } Braces	Enclose an optional value or a list of optional arguments. You can specify one or more values or arguments. For example, in the syntax `reboot {<date> <time> | cancel}` you can specify either a particular date and time combination, or the keyword `cancel` to cancel a scheduled reboot. If you do not specify a value, the system prompt asks if you want to reboot the routing switch now. Do not type the braces.

Line-Editing Keys

Table 5.2 describes the line-editing keys available using the CLI.

Table 5.2: Line-Editing Keys

Key(s)	Description
Backspace	Deletes characters to the left of the cursor and shifts the remainder of the line to the left.
Delete or Ctrl + D	Deletes character at the cursor position and shifts the remainder of line to the left.
Ctrl + K	Deletes characters from the cursor position to the end of the line.
Ctrl + U	Deletes characters from the cursor to the beginning of the line.
Ctrl + W	Deletes the previous word.
Left Arrow	Moves the cursor to the left.
Right Arrow	Moves the cursor to the right.
Home or Ctrl + A	Moves the cursor to first character on the line.
End or Ctrl + E	Moves the cursor to last character on the line.
Ctrl + L	Clears the screen and moves the cursor to the beginning of the line.
Up Arrow or Ctrl + P	Displays the previous command in the command history buffer and places the cursor at the end of the command.
Down Arrow or Ctrl + N	Displays the next command in the command history buffer and places the cursor at the end of the command.

Command History

The local management software stores the last 49 commands you entered. You can display a list of these commands by using this command:

```
history
```

Common Commands

Table 5.3 describes common commands used to manage the 480T routing switch. Commands specific to particular features are described in detail throughout the guide. For detailed command information use the Quick Reference Guide that accompanies this user manual. Press the Tab key in the command line interface for more command options.

Table 5.3: Common Commands

Command	Description
clear session <number>	Terminates a Telnet session from the switch.
configure account <username> {<password>}	Configures a user account password. Passwords can have no characters up to a maximum of 32 characters. User names and passwords are case-sensitive.
configure banner	Configures the banner string. You can enter up to 24 rows of 79-column text that is displayed before the login prompt of each session. To terminate the command, apply the banner then press Enter at the beginning of a line . To clear the banner, press Enter at the beginning of the first line.
configure ports [all | mgmt | <portlist>] auto off {speed [100 | 1000]} duplex [half | full]	Manually configures Ethernet port speed and duplex setting of one or more ports on a switch.

Table 5.3: Common Commands (continued)

Command	Description
configure time <date> <time>	Configures the system date and time. The format is as follows: `mm/dd/yyyy hh:mm:ss` The time uses a 24-hour clock format.
configure timezone <gmt_offset> {autodst | noautodst}	Configures the time zone information to the configured offset from Greenwich Mean Time (GMT) time. The format of `gmt_offset` is +/- minutes from GMT time. Specify: `autodst`—Enables automatic daylight saving time change. `noautodst`—Disables automatic daylight saving time change. The default setting is `autodst`.
configure vlan <name> ipaddress <ip_address> {<mask>}	Configures an IP address and subnet mask for a VLAN.
create account [admin | user] <username> {encrypted} {<password>}	Creates a user account. The command is available to admin-level users and users with RADIUS§ command authorization. The username can be between 1 and 32 characters. The password can be between 0 and 32 characters.
create vlan <name>	Creates a VLAN.
delete account <username>	Deletes a user account.
delete vlan <name>	Deletes a VLAN.
disable bootp vlan [<name> | all]	Disables BOOTP for one or more VLANs.
disable cli-config-logging	Disables logging of CLI commands to the Syslog.

Table 5.3: Common Commands (continued)

Command	Description
disable clipaging	Disables pausing of the screen display when a **show** command output reaches the end of the page.
disable idletimeout	Disables the timer that disconnects all sessions. Once disabled, console sessions remain open until the switch is rebooted or you log off. Telnet sessions remain open until you close the Telnet client.
disable port [all | mgmt | <portlist>]	Disables a port on the switch.
disable telnet	Disables Telnet access to the switch.
disable web	Disables Web access to the switch.
enable bootp vlan [<name> | all]	Enables BOOTP for one or more VLANs.
enable cli-config-logging	Enables the logging of CLI configuration commands to the Syslog for auditing purposes. The default setting is enabled.
enable clipaging	Enables pausing of the screen display when **show** command output reaches the end of the page. The default setting is enabled.
enable idletimeout	Enables a timer that disconnects all sessions (both Telnet and console) after **20** (in minutes) of inactivity. The default setting is disabled.
enable license full_L3 <license_key>	Enables a particular software feature license. Specify <**license_key**> as an integer. The command **unconfigure switch all** does not clear licensing information. This license cannot be disabled after it is enabled on the switch.

Table 5.3: Common Commands (continued)

Command	**Description**
enable telnet {access-profile [<access_profile> | none]} {port <tcp_port_number>}	Enables Telnet access to the switch. By default, Telnet is enabled with no access profile, and uses Transmission Control Protocol (TCP) port number 23. To cancel a previously configured access profile, use the **`none`** option.
enable web {access-profile [<access_profile> | none]} {port <tcp_port_number>}	Enables Web access to the switch. By default, Web access is enabled with no access profile, using TCP port number 80. Use the **`none`** option to cancel a previously configured access profile. Reboot the switch for this command to take effect.
history	Displays the previous 49 commands entered on the switch.
show banner	Displays the user-configured banner.
unconfigure switch {all}	The **`unconfigure switch`** command resets parameters to factory defaults, except defined user accounts, and date and time information. To reset user accounts and date and time, specify the keyword **`all`** which erases the selected configuration image in flash memory and reboots.

Configuring Management Access

The local management software supports these two levels of management:

- User
- Administrator

In addition to these management levels, you can optionally use an external RADIUS server to provide CLI command authorization checking for each command.

For more information on RADIUS, refer to "RADIUS Client" on page 66."

A user-level account has viewing access to all manageable parameters, with the exception of these:

- User account database
- SNMP community strings

User Account

With a user-level account you can use the **`ping`** command to test device connectivity, and change the password assigned to the account name. When you log on the command-line prompt ends with a (>) sign. For example:

```
switch480T:2>
```

Administrator Account

Using an administrator-level account, you can view and change all routing switch parameters. You can also add and delete users, and change the password associated with any account name.

As an administrator you can also disconnect a management session connected through Telnet. If this happens, the user logged on through the Telnet connection is notified that the session was terminated.

When you log on with administrator capabilities, the command-line prompt ends with a (#) sign. For example:

```
switch480T:18#
```

Prompt Text

The prompt text is taken from the SNMP **`sysname`** setting (see Table 5.8, “SNMP Configuration Commands,” on page 64). The number that follows the colon indicates the sequential line/command number.

If an asterisk (*) appears in front of the command-line prompt, it indicates that you have configuration changes that have not been saved. For example:

```
*switch480T:19#
```

Default Accounts

The switch is configured with two default accounts. as shown in Table 5.4.

Table 5.4: Default Accounts

Account Name	Access Level
admin	This user can access and change all manageable parameters. The admin account cannot be deleted.
user	This user can view (but not change) almost all manageable parameters. However, this user cannot view the user account database or the SNMP community strings.

Changing the Default Password

Default accounts do not have passwords assigned to them. User-assigned passwords must be between 0 and 32 characters.

Passwords are case sensitive.

To add a password to the default admin account:

1. Log in to the switch using the name ***admin***.
2. At the password prompt, press Enter.
3. Enter this command:

   ```
   configure account admin
   ```

4. Enter the new password at the prompt.
5. Re-enter the password for verification.

To add a password to the default user account:

1. Log in to the switch using the name ***admin***.
2. At the password prompt, press Enter, or enter the password that you have configured for the admin account.
3. Add a default user password using this command:

```
configure account user
```

4. Enter the new password at the prompt.
5. Re-enter the new password at the prompt.

Creating a Management Account

If you forget your password while logged out of the command-line interface, contact your local technical support representative.

The 480T routing switch can have a total of 16 management accounts. You can use the default names (admin and user), or you can create new names and passwords for the accounts. Account passwords can be between 0 and 32 characters. Do not use Ctrl + key or Alt + key.

To create a management account:

1. Log in to the switch as ***admin***.
2. At the password prompt, press Enter, or enter the password that you have configured for the admin account.
3. Add a new user account with this command:

```
create account [admin | user] <username>
```

4. Enter the password at the prompt.
5. Re-enter the password for verification.

Viewing Accounts

To view the accounts you have created, you must have administrator privileges. Use this command to see the accounts:

```
show accounts
```

Deleting an Account

To delete an account, you must have administrator privileges. Use this command to delete an account:

```
delete account <username>
```

The account name admin cannot be deleted.

Domain Name Service Client

The Domain Name Service (DNS) client augments these commands, to allow them to accept either IP addresses or host names:

- **`telnet`**
- **`download [bootrom | configuration | image]`**
- **`upload configuration`**
- **`ping`**
- **`traceroute`**

Also, you can use the **`nslookup`** utility to return the IP address of a host name.

Table 5.5 describes the commands used to configure DNS. Press the Tab key in the command line interface for more command options.

Table 5.5: DNS Commands

Command	Description
configure dns-client add <ipaddress>	Adds a DNS name server(s) to the available server list for the DNS client. You can configure up to three name servers.
configure dns-client default-domain <domain_name>	Configures the domain that the DNS client uses if a fully qualified domain name is not entered. For example, if the default domain is configured to be **`intel.com,`** executing **`ping support`** searches for support@intel.com.
configure dns-client delete <ipaddress>	Removes a DNS server.
nslookup <hostname>	Displays the IP address of the requested host.
show dns-client	Displays the DNS configuration.

Real-time Basic Connectivity Checking

Use these commands to check basic connectivity:

- **ping**
- **traceroute**

Ping

You can use the **ping** command to send Internet Control Message Protocol (ICMP) echo messages to a remote IP device. The **ping** command is available for both the user and administrator privilege level.

The **ping** command syntax is:

```
ping {continuous} {start-size <start_size> {- end-size <end_size>}} [<ip_address> | <hostname>] {from <src_address>} {with record-route}
```

Options for the ping command are described in Table 5.6. Press the Tab key in the command line interface for more command options.

Table 5.6: Ping Command Parameters

Parameter	Description
continuous	Specifies Internet Control Message Protocol (ICMP) echo messages to be sent continuously. To interrupt this option, press any key.
size <n>	Specifies the size of the ICMP request. If both **start-size** and **end-size** are specified, ICMP requests are transmitted using increments of 1 byte per packet. If no **end-size** is specified, packets of **start-size** are sent.
<ipaddress>	Specifies the IP address of the host.
<hostname>	Specifies the name of the host. To use the **hostname**, first configure DNS.

Table 5.6: Ping Command Parameters (continued)

Parameter	Description
`from`	Uses the specified source address in the ICMP packet. If not specified, the address of the transmitting interface is used.
`with record-route`	Decodes the list of recorded routes and displays them when the ICMP echo reply is received.

Traceroute

The **`traceroute`** command enables you to trace the routed path between the switch and a destination endstation. The **`traceroute`** command syntax is:

```
traceroute [<ip_address> | <hostname>] {from
<src_ipaddress>} {ttl <TTL>} {port <port>}
```

where:

- `ip_address` is the IP address of the destination endstation.
- `hostname` is the host name of the destination endstation. To use the host name, first configure DNS.
- `from` uses the specified source address in the ICMP packet. If not specified, the address of the transmitting interface is used.
- `ttl` configures the switch to trace up to the time-to-live number of the switch.
- `port` uses the specified UDP port number.

Methods of Managing the Switch

See "Using Intel® Device View" on page 23.

You can manage the switch by either connecting a terminal (or workstation with terminal-emulation software) to the console port to access the CLI or by using TCP/IP through one of the switch ports or through the dedicated 10/100 Mbps unshielded twisted pair (UTP) Ethernet management port to access the switch remotely.

You can use Telnet, a Web browser, or an SNMP manager to manage the switch remotely. There can be one console session, one Web session or eight concurrent Telnet sessions.

Using the Console Interface

You can access the built-in CLI of the 480T routing switch through the 9-pin RS-232 port located on the back of the switch.

After the connection is established, the switch prompt appears, so you can log in.

Using the 10/100 UTP Management Port

The 480T routing switch has a dedicated 10/100 Mbps UTP management port. This port provides dedicated remote access to the switch using TCP/IP. It supports these management methods:

- Telnet using the CLI interface
- Intel Device View access using a Web browser
- SNMP access using SNMP manager

The management port is a DTE port, and is not capable of supporting switching or routing functions. The TCP/IP configuration for the management port is done using the same syntax as used for VLAN configuration. The VLAN ***mgmt*** comes pre-configured with only the 10/100 Mbps management port as a member.

You can configure the IP address, subnet mask, and default router for the VLAN ***mgmt***, using these commands:

- **`configure vlan mgmt ipaddress <ip_address>/ <subnet_mask>`**
- **`configure iproute add default <gateway>`**

Using Telnet

Most workstations with a Telnet facility can communicate with the 480T routing switch over a TCP/IP network.

Up to eight active Telnet sessions can access the switch concurrently. If **`idletimeouts`** are enabled, the Telnet connection will time out after 20 minutes of inactivity. If a connection to a

Telnet session is lost inadvertently, the switch terminates the session within two hours.

Before you can start a Telnet session, you must set up the IP parameters described in the section "Configuring Switch IP Parameters" on page 55.. Telnet is enabled by default.

To open the Telnet session, you must specify the IP address of the device that you want to manage. Check the user manual supplied with the Telnet facility if you are unsure of how to do this.

After the connection is established, you will see the switch prompt and you can log in.

Connecting to Another Host Using Telnet

Use this command to Telnet from the current CLI session to another host:

```
telnet [<ipaddress> | <hostname>] {<port_number>}
```

If the TCP port number is not specified, the Telnet session defaults to port 23. Only VT100 emulation is supported.

Configuring Switch IP Parameters

To manage the routing switch through Telnet connection or by using an SNMP Network Manager, you must first configure the switch IP parameters.

Using a BOOTP Server

If you are using IP and you have a Bootstrap Protocol (BOOTP) server set up correctly on your network, you must add the following information to the BOOTP server:

- Media Access Control (MAC) address found on the rear label of the switch (or use the **`show switch`** command)
- IP address
- Subnet address mask (optional)

Find the switch's MAC address on the rear label of the switch.

After this is done, the IP address and subnet mask for the routing switch is downloaded automatically. You can then start managing the switch without further configuration.

You can enable BOOTP on a per-VLAN basis using this command:

```
enable bootp vlan [<name> | all]
```

By default, BOOTP is enabled on the ***default*** VLAN.

If you configure the 480T routing switch to use BOOTP, the switch IP address is not retained through a power cycle, even if the configuration is saved. To retain the IP address through a power cycle, you must configure the IP address of the VLAN using the command-line interface, Telnet, or Web interface.

All VLANs within a switch that are configured to use BOOTP to get their IP address use the same MAC address. Therefore, if you are using BOOTP relay through a router, the BOOTP server must be capable of differentiating its relay based on the gateway portion of the BOOTP packet.

Manually Configuring the IP Settings

For more information on DHCP/BOOTP relay, refer to "IP Unicast Routing" on page 189.

If you are using IP without a BOOTP server, you must enter the IP parameters for the switch in order for the SNMP Network Manager, Telnet software, or Web interface to communicate with the device.

IP addresses are always assigned to a VLAN. You can assign multiple IP addresses to the switch.

To assign IP parameters to the switch:

1. Log in to the switch with administrator privileges.
2. Assign an IP address and subnet mask to a VLAN.

The switch comes configured with a default VLAN named ***default***. To use Telnet or an SNMP Network Manager, you must have at least one VLAN on the switch, and it must be assigned an IP address and subnet mask.

For information on creating and configuring VLANs, see "Virtual LANs (VLANs)" on page 95.

To manually configure the IP settings:

1. Connect a terminal or workstation running terminal-emulation software to the console port.
2. At your terminal, press Enter one or more times until you see the login prompt.
3. If you are logging in for the first time, use the default user name ***admin*** to log in with administrator privileges. For example:

```
login: admin
```

Administrator capabilities enable you to access all switch functions. The default user names have no passwords assigned.

4. If you have been assigned a user name and password with administrator privileges, enter them at the login prompt and press Enter.

   When you have successfully logged in, the command-line prompt displays the name of the switch.

5. Assign an IP address and subnetwork mask for the default VLAN using this command:

   ```
   configure vlan <name> ipaddress <ipaddress>
   {<subnet_mask>}
   ```

   For example:

   ```
   configure vlan default ipaddress 123.45.67.8
   255.255.255.0
   ```

   Your changes take effect immediately.

   Generally, when configuring any IP addresses for the switch, you can express a subnet mask using dotted decimal notation, or classless inter-domain routing notation (CIDR).

   CIDR uses a forward slash plus the number of bits in the subnet mask. Using CIDR notation, the command identical to the one above would be:

   ```
   configure vlan default ipaddress 123.45.67.8/24
   ```

6. Configure the default route for the switch using this command:

   ```
   configure iproute add default <gateway>
   {<metric>}
   ```

   For example:

   ```
   configure iproute add default 123.45.67.1
   ```

7. Save your configuration changes so that they are in effect after the next switch reboot, using this command.

   ```
   save
   ```

8. Log out of the switch using the command:

   `logout` or `quit`

Disconnecting a Telnet Session

An administrator-level account can disconnect a management session that is established through Telnet connection. If this happens, the user logged in through Telnet is notified that the session is terminated.

To terminate a Telnet session:

1. Log in to the switch with administrator privileges.
2. Determine the session number of the session you want to terminate by using this command:

   ```
   show session
   ```

3. Terminate the session by using this command:

   ```
   clear session <session_number>
   ```

Controlling Telnet Access

See "Using Access Profiles" on page 59.

By default, Telnet services are enabled on the routing switch. You can restrict Telnet access using an access profile. An access profile permits or denies a named list of IP addresses and subnet masks. To configure Telnet to use an access profile, use this command:

```
enable telnet {access-profile [<access_profile> |
none]} {port <tcp_port_number>}
```

Use the **none** option to remove a previously configured access profile.

To display the status of Telnet, use this command:

```
show management
```

You must be logged in as an administrator to enable or disable Telnet.

To disable Telnet, use this command:

```
disable telnet
```

To re-enable Telnet on the switch, use this command at the console port:

```
enable telnet
```

Using Access Profiles

An access profile permits or denies a named list of IP addresses and subnet masks. To use access profiles, first define the list, and then apply the named list to the desired application.

Access profiles are used by several routing switch features as a way to restrict access. Applications that use access profiles for remotely managing the switch are:

- SNMP read-only access
- SNMP read-write access
- Telnet
- Web access

See "Access Policies" on page 309.

Access profiles can also be used in association with access policies that control the flow of traffic.

Creating an Access Profile

Do not confuse access profiles with access policies.

You can use access profiles to specifically permit or deny users access to an application. You restrict access by assigning an access profile to the service that is being used for remote access.

When you create and name an access profile to restrict access to a certain application, you then need to configure the application to use the named access profile. Otherwise, no restrictions are applied.

Use the commands listed in Table 5.7 to create and configure access profiles. For further access profile commands refer to Table 17.3 on page 335. Press the Tab key in the command line interface for more command options.

Table 5.7: Access Profile Configuration Commands

Command	Description
configure access-profile <access_profile> add {vlan <name> | ipaddress <ipaddress> <mask>}	Adds an IP address or VLAN name to the access profile. The entry must be of the same type as the access profile (for example, IP address).

Table 5.7: Access Profile Configuration Commands (continued)

Command	Description
configure access-profile <access_profile> delete {vlan <name> | ipaddress <ipaddress> <mask>}	Deletes an IP address or VLAN name from the access profile.
configure access-profile <access_profile> mode [permit | deny | none]	Configures the access profile to one of the following: **`permit`**—Allows the addresses that match the access profile description. **`deny`**—Denies the addresses that match the access profile description. The default setting is **`permit`**.
create access-profile <access_profile> type [as-path] [bgp-community] ipaddress | ipxret | ipxnode | ipxsap	Creates an access profile. After the access profile is created, you can add one or more addresses to it, and you can use the profile to control a specific routing protocol.
delete access-profile <access_profile>	Deletes an access profile.
show access-profile <access_profile>	Displays access profile related information for the switch.

The subnet mask specified in the access profile command is interpreted as a reverse mask. A reverse mask indicates the bits that are significant in the IP address and specifies the part of the address that must match the IP address to which the profile is applied.

If you configure an IP address as an exact match to be specifically denied or permitted, use a mask of /32 (for example, 141.251.24.28/32).

If the IP address represents a subnet address that you want to deny or permit, then configure the mask to cover only the subnet portion (for example, 141.251.10.0/24).

If you are using classless subnet masking (CIDR), the same logic applies, but the configuration is more complex. For example, the address 141.251.24.128/27 represents any host from subnet 141.251.24.128.

Access Profile Rules

These rules apply when using access profiles:

- Only one access profile can be applied to each application.
- The access profile can either permit or deny the entries in the profile.
- The same access profile can be applied to more than one application.

Access Profile Example

The following example creates an access profile named ***testpro***, and denies access for the device with the IP address 192.168.10.10:

```
create access-profile testpro type ipaddress
configure access-profile testpro mode deny
configure access-profile testpro add ipaddress
192.168.10.10/32
```

The following command applies the access profile ***testpro*** to Telnet:

```
enable telnet access-profile testpro
```

To view the contents of an access profile, use this command:

```
show access-profile <access_profile>
```

To view the Telnet configuration, use this command:

```
show management
```

Using Web Device Manager

The Intel Web Device Manager is device-management software running in the routing switch that enables you to access the switch over a TCP/IP network using a Web browser.

For more information, refer to "Using Web Device Manager" on page 33.

You should use a Web browser that supports frames (such as Netscape Navigator§ 3.0 or later, or Microsoft Internet Explorer§ 3.0 or later) to manage the switch over a TCP/IP network.

Access the default home page of the switch using this command:

```
http://<ipaddress>
```

When you access the home page of the switch the Logon screen appears.

Controlling Web Access

By default, Web access is enabled on the routing switch. You can restrict access through the Web Device Manager using an access profile, which permits or denies access to a named list of IP addresses and subnet masks.

For more information on assigning an IP address, refer to "Configuring Switch IP Parameters" on page 55.

You can configure Web access to use an access profile using this command:

```
enable web {access-profile <access-profile> | none}
{port <tcp_port_number>}
```

Use the **none** option to remove a previously configured access profile.

To display the status of Web access, use this command:

```
show management
```

To disable Web access, use this command:

```
disable web
```

To re-enable Web access, use this command:

```
enable web {access-profile <access-profile> | none}
{port <tcp_port_number>}
```

When you disable or enable Web Device Manager, you must reboot the switch for the changes to take effect. Apply an access profile only when Web Device Manager is enabled.

Simple Network Management Protocol (SNMP)

Any network manager running the Simple Network Management Protocol (SNMP) can manage the 480T routing switch, provided the Management Information Base (MIB) feature of the 480T routing switch is installed correctly on the management station. Each Network Manager provides its own user interface to the management facilities.

Accessing Switch Agents

To have access to the SNMP agent in the routing switch, at least one VLAN must have an IP address assigned to it.

For more information on assigning IP addresses, refer to Table 5.3 on page 44.

Supported MIBs

Along with private MIBs, the routing switch supports the MIBs listed in "Technical Specifications and Supported Limits" on page 431.

Configuring SNMP Settings

You can configure the following SNMP parameters on the routing switch:

- **Authorized trap receivers**—An authorized trap receiver can be one or more network management stations on your network. The switch sends SNMP traps to all trap receivers. You can have a maximum of 16 trap receivers configured for each switch. .
- **SNMP read access**—The ability to read SNMP information can be restricted through the use of an access profile. An access profile permits or denies a named list of IP addresses and subnet masks.

  To configure SNMP read access to use an access profile, use the command:

  ```
  configure snmp access-profile readonly
  [<access_profile> | none]
  ```

  Use the `none` option to remove a previously configured access profile.
- **SNMP read/write access**—The ability to read and write SNMP information can be restricted through the use of an access profile. An access profile permits or denies a named list of IP addresses and subnet masks.

  To configure SNMP read/write access to use an access profile, use the command:

  ```
  configure snmp access-profile readwrite
  [<access_profile> | none]
  ```

  Use the `none` option to remove a previously configured access profile.

- **Community strings**—Allows a simple method of authentication between the 480T routing switch and the remote Network Manager. There are two types of community strings on the switch. Read community strings provide read-only access to the switch. The default read-only community string is public. Read-write community strings provide read and write access to the switch. The default read-write community string is private. A total of eight community strings can be configured on the switch. The community string for all authorized trap receivers must be configured on the switch for the trap receiver to receive switch-generated traps. SNMP community strings can contain up to 127 characters.
- **System contact** (optional)—A text field where you can enter the name of the person(s) responsible for managing the switch.
- **System name**—The name you have assigned to this switch. The default name is switch480T.
- **System location** (optional)—Use this to enter an optional location for this switch.

Table 5.8 describes SNMP configuration commands. Press the Tab key in the command line interface for more command options.

Table 5.8: SNMP Configuration Commands

Command	**Description**
configure snmp access-profile readonly [<access_profile> | none]	Assigns an access profile that limits which stations have read-only access to the switch.
configure snmp access-profile readwrite [<access_profile> | none]	Assigns an access profile that limits which stations have read-write access to the switch.
configure snmp add trapreceiver <ipaddress> community <string>	Adds the IP address of a specified trap receiver. The IP address can be a unicast, multicast, or broadcast address. A maximum of 16 trap receivers is allowed.

Table 5.8: SNMP Configuration Commands (continued)

Command	Description
configure snmp community [readonly | readwrite] {encrypted} <string>	Adds an SNMP read or read/write community string. The default `readonly` community string is `public`. The default `readwrite` community string is `private`. Each community string can have a maximum of 127 characters, and can be enclosed by double quotation marks.
configure snmp delete trapreceiver [<ip_address> community <string> | all]	Deletes the IP address of a specified trap receiver or all authorized trap receivers.
configure snmp syscontact <string>	Configures the name of the system contact. A maximum of 255 characters is allowed.
configure snmp syslocation <string>	Configures the location of the switch. A maximum of 255 characters is allowed.
configure snmp sysname <string>	Configures the name of the switch. A maximum of 32 characters is allowed. The default sysname is the model name of the device (for example, `switch480T)`. The `sysname` appears in the switch prompt.
disable snmp access	Disables SNMP access on the switch. Disabling SNMP access does not affect the SNMP configuration (for example, community strings).
disable snmp traps	Prevents SNMP traps from being sent from the switch. This does not clear the SNMP trap receivers that have been configured.
enable snmp access	Enables SNMP support.
enable snmp traps	Enables SNMP trap support.
unconfigure management	Restores default values to all SNMP-related entries.

Displaying SNMP Settings

To display the SNMP settings configured on the routing switch, use this command:

```
show management
```

This command displays the following information:

- Enable/disable state for Telnet, SNMP, and Web access, along with access profile information
- SNMP community strings
- Authorized SNMP station list
- SNMP trap receiver list
- RMON polling configuration
- Login statistics

SNMP enhancements allow the ifMIB to display the port number for physical ports and VLAN name for the VLANs index.

Authenticating Users

The routing switch uses two methods to authenticate users who login to the switch:

- RADIUS§ client
- TACACS+ (Terminal Access Controller Access Control System Plus)

RADIUS Client

Remote Authentication Dial In User Service (RADIUS, RFC 2138) allows you to authenticate and centrally administer access to network nodes. The 480T routing switch RADIUS client implementation enables authentication for Telnet, Web interface, or console access to the switch.

You cannot configure RADIUS and TACACS+ at the same time.

You can define a primary and secondary RADIUS® server for the routing switch to contact.

When a user attempts to log on to the switch using Telnet, HTTP, or the console, the request is relayed to the primary RADIUS server, and then to the secondary RADIUS server, if the primary does not respond.

If the RADIUS client is enabled, but access to the RADIUS primary and secondary servers fail, the routing switch uses its local database for authentication.

The privileges assigned to the user (admin versus non-admin) at the RADIUS server take precedence over the configuration in the local switch database.

Per-Command Authentication Using RADIUS

Use RADIUS to perform per-command authentication. Per-command authentication allows you to define several levels of user capabilities that determine which set of commands the user has access to based on the RADIUS username and password.

There is no need to configure any additional switch parameters to take advantage of this capability. The RADIUS server implementation automatically negotiates the per-command authentication capability with the switch.

Configuring RADIUS Client

You can define primary and secondary server communication information. Also for each RADIUS server, you can specify the RADIUS port number to use when talking to the RADIUS server. The default port value is 1645. The client IP address is the IP address used by the RADIUS server for communicating with the 480T routing switch.

RADIUS commands are described in Table 5.9. Press the Tab key in the command line interface for more command options.

Table 5.9: RADIUS® Commands

Command	Description
configure radius [primary | secondary] server [<ipaddress> | <hostname>] {<udp_port>} client-ip <ipaddress>	Configures the primary and secondary RADIUS§ server. Specify the following: • `[primary | secondary]`—Either the primary or secondary RADIUS server. • `[<ipaddress> | <hostname>]`—The IP address or host name of the server being configured. • `<udp_port>`—The UDP port to use to contact the RADIUS server. The default UDP port setting is 1645. • `client-ip <ipaddress>`–The IP address used by the switch to identify itself when communicating with the RADIUS server. The RADIUS server defined by this command is used for user-name authentication and CLI command authentication.
configure radius [primary | secondary] shared-secret {encrypted} <string>	Configures the authentication string used to communicate with the RADIUS server.
configure radius-accounting [primary | secondary] shared-secret {encrypted} <string>	Configures the authentication string used to communicate with the RADIUS accounting server.
disable radius	Disables the RADIUS client.
disable radius-accounting	Disables RADIUS accounting.

Table 5.9: RADIUS® Commands (continued)

Command	Description
configure radius-accounting [primary | secondary] server [<ipaddress> | <hostname>] {<udp_port>} client-ip <ipaddress>	Configures the RADIUS accounting server. Specify the following: • `[primary	secondary]`—Either the primary or secondary RADIUS server. • `[<ipadress>	<hostname>]`—The IP address or host name of the server being configured. • `<udp_port>`—The UDP port to use to contact the RADIUS server. The default UDP port setting is 1646. • `client-ip <ipaddress>`–The IP address used by the switch to identify itself when communicating with the RADIUS server. The accounting server and the RADIUS authentication server can be the same.
enable radius	Enables the RADIUS client. When enabled, all Web and CLI logins are sent to the RADIUS servers for authentication. When used with a RADIUS server that supports routing switch CLI authorization, each CLI command is sent to the RADIUS server for authentication before it is executed.
enable radius-accounting	Enables RADIUS accounting. The RADIUS client must also be enabled.
show radius	Displays the current RADIUS and RADIUS accounting client configuration and statistics.
show radius-accounting	Displays the current RADIUS accounting client configuration and statistics.

RADIUS RFC 2138 Attributes

The RADIUS RFC 2138 optional attributes supported are:

- User-Name
- User-Password
- Service-Type
- Login-IP-Host

Configuring TACACS+

Terminal Access Controller Access Control System Plus (TACACS+) is a means for providing authentication, authorization, and accounting on a centralized server, similar in function to a RADIUS client.

The routing switch version of TACACS+ is used to authenticate prospective users who are attempting to administer the switch. TACACS+ is used to communicate between the switch and an authentication database.

You cannot use TACACS+ and RADIUS at the same time.

You can configure two TACACS+ servers, specifying the primary server address, secondary server address, and UDP port number to be used for TACACS+ sessions.

Table 5.10 describes the commands that are used to configure TACACS+. Press the Tab key in the command line interface for more command options.

Table 5.10: TACACS+ Commands

Command	Description
configure tacacs [primary | secondary] server [<ipaddress> | <hostname>] {<udp_port>} client-ip <ipaddress>	Configures the server information for a TACACS+ server. Specify the following: • `primary	secondary`—Specifies primary or secondary server configuration. To remove a server, use the address `0.0.0.0.` • `<ipaddress>	<hostname>`—The IP address or hostname of the TACACS+ server. • `<udp_port>`—Optionally specifies the UDP port to be used. • `client-ip`—Specifies the IP address used by the switch to identify itself when communicating with the TACACS+ server.
configure tacacs [primary | secondary] shared-secret {encrypted} <string>	Configures the shared secret string used to communicate with the TACACS+ server.
configure tacacs-accounting [primary | secondary] server [<ipaddress> | <hostname>] {<udp_port>} client-ip <ipaddress>	Configures the TACACS+ accounting server. You can use the same server for accounting and authentication.
configure tacacs-accounting [primary | secondary] shared-secret {encrypted} <string>	Configures the shared secret string used to communicate with the TACACS+ accounting server.
disable tacacs	Disables TACACS+.
disable tacacs-accounting	Disables TACACS+ accounting.
disable tacacs-authorization	Disables CLI command authorization.
enable tacacs	Enables TACACS+. Once enabled, all Web and CLI logins are sent to one of the two TACACS+ servers for login name authentication and accounting.

Table 5.10: TACACS+ Commands (continued)

Command	Description
enable tacacs-accounting	Enables TACACS+ accounting. If accounting is used, the TACACS+ client must also be enabled.
enable tacacs-authorization	Enables CLI command authorization. When enabled, each command is transmitted to the remote TACACS+ server for authorization before the command is executed.
show tacacs	Displays the current TACACS+ configuration and statistics.
show tacacs-accounting	Displays the current TACACS+ accounting client configuration and statistics.
unconfigure tacacs {server [primary | secondary]}	Unconfigures the TACACS+ client configuration.
unconfigure tacacs-accounting {server [primary | secondary]}	Unconfigures the TACACS+ accounting client configuration.

Simple Network Time Protocol (SNTP)

Therouting switch supports the client portion of the Simple Network Time Protocol (SNTP) Version 3 based on RFC1769. The switch can use SNTP to update and synchronize its internal clock from a Network Time Protocol (NTP) server.

When SNTP is enabled, the switch sends out a periodic query to the indicated NTP server, or the switch listens to broadcast NTP updates. The routing switch also supports the configured setting for Greenwich Mean time (GMT) offset and the use of daylight saving time.

Configuring and Using SNTP

To use SNTP:

1 Identify the host(s) that are configured as NTP server(s).

2 Identify the preferred method for obtaining NTP updates.

The options are for the NTP server to send out broadcasts, or for switches using NTP to query the NTP server(s) directly. A combination of both methods is possible.

3 Configure the Greenwich Mean Time (GMT) offset and daylight saving time preference. NTP updates are distributed using GMT time.

To properly display the local time in logs and other timestamp information, the switch should be configured with the appropriate offset to GMT based on geographical location. Table 5.11 describes GMT offsets.

The command syntax to configure GMT offset and usage of daylight saving time is as follows:

```
configure timezone <GMT_offset> {autodst |
noautodst}
```

The GMT_OFFSET is in +/- minutes from the GMT time. You can enable or disable Automatic daylight saving time (DST) changes. The default setting is enabled.

4 Enable the SNTP client using this command:

```
enable sntp-client
```

Once enabled, the switch sends out a periodic query to the NTP servers (if configured) or listens to broadcast NTP updates from the network. The network time information is automatically saved in the on-board real-time clock.

5 If you would like this switch to use a directed query to the NTP server, configure the switch to use the NTP server(s). If the switch listens to NTP broadcasts, skip this step. To configure the 480T routing switch to use a directed query, use this command:

```
configure sntp-client [primary | secondary]
server [<ip_address> | <hostname>]
```

NTP queries are first sent to the primary server. If the primary server does not respond within one second, or if it is not synchronized, the switch queries the secondary server (if configured).

If the switch cannot obtain the time, it restarts the query process. Otherwise, the switch waits for the `sntp-client update interval` before querying again.

6 Optionally, you can change the interval for which the SNTP client updates the real-time clock of the switch using this command:

```
configure sntp-client update-interval <seconds>
```

The default `sntp-client update-interval` value is `64`

7 You can verify the configuration using these commands:

```
show sntp-client
```

8 This command provides configuration and statistics associated with SNTP and its connectivity to the NTP server:

```
show switch
```

This command indicates the GMT offset, daylight saving time, and the current local time.

Table 5.11: Greenwich Mean Time Offsets

GMT Offset in Hours	GMT Offset in Minutes	Common Time Zone References	Geographical Reference
+0:00	+0	GMT - Greenwich Mean UT or UTC - Universal (Coordinated) WET - Western European	London, England; Dublin, Ireland; Edinburgh, Scotland; Lisbon, Portugal; Reykjavik, Iceland; Casablanca, Morocco
-1:00	-60	WAT - West Africa	Cape Verde Islands
-2:00	-120	AT - Azores	Mid-Atlantic
-3:00	-180		Brasilia, Brazil; Buenos Aires, Argentina; Georgetown, Guyana;
-4:00	-240	AST - Atlantic Standard	Caracas, La Paz
-5:00	-300	EST - Eastern Standard	Bogota, Columbia; Lima, Peru; New York, NY, USA;

Table 5.11: Greenwich Mean Time Offsets (continued)

GMT Offset in Hours	GMT Offset in Minutes	Common Time Zone References	Geographical Reference
-6:00	-360	CST - Central Standard	Chicago, Illinois, USA; Mexico City, Mexico; Saskatchewan, Canada
-7:00	-420	MST - Mountain Standard	Salt Lake City, Utah, USA; Alberta, Canada
-8:00	-480	PST - Pacific Standard	Los Angeles, CA. USA; Seattle, WA, USA
-9:00	-540	YST - Yukon Standard	Whitehorse, Alaska, USA; Yukon Territory, Canada
-10:00	-600	AHST - Alaska-Hawaii Standard CAT - Central Alaska HST - Hawaii Standard	Honolulu, Hawaii
-11:00	-660	NT - Nome	Midway Islands, Samoa
-12:00	-720	IDLW - International Date Line West	Eniwitok, Kwajalein
+1:00	+60	CET - Central European FWT - French Winter MET - Middle European MEWT - Middle European Winter SWT - Swedish Winter	Paris, France; Berlin, Germany; Amsterdam, The Netherlands; Brussels, Belgium; Vienna, Austria; Madrid, Spain; Rome, Italy; Bern, Switzerland; Stockholm, Sweden; Oslo, Norway
+2:00	+120	EET - Eastern European, Russia Zone 1	Athens, Greece; Helsinki, Finland; Istanbul, Turkey; Jerusalem, Israel; Harare, Zimbabwe
+3:00	+180	BT - Baghdad, Russia Zone 2	Kuwait; Nairobi, Kenya; Riyadh, Saudi Arabia; Moscow, Russia; Tehran, Iran

Table 5.11: Greenwich Mean Time Offsets (continued)

GMT Offset in Hours	GMT Offset in Minutes	Common Time Zone References	Geographical Reference
+4:00	+240	ZP4 - Russia Zone 3	Abu Dhabi, UAE; Muscat; Tblisi; Volgograd; Kabul
+5:00	+300	ZP5 - Russia Zone 4	Islamabad, Karachi, Tashkent, Russia
+5:30	+330	IST – India Standard Time	Bombay, Calcutta, Madras, New Delhi, Pune, Allahabad, India
+6:00	+360	ZP6	Dhaka, Almaty
+7:00	+420	WAST - West Australian Standard	Bangkok, Jakarta
+8:00	+480	CCT - China Coast, Russia Zone 7	Beijing, Hong Kong, Perth, Singapore, Taipei
+9:00	+540	JST - Japan Standard, Russia Zone 8	Tokyo, Japan; Osaka, Sapporo Seoul, Yakutsk
+10:00	+600	EAST - East Australian Standard GST - Guam Standard Russia Zone 9	Brisbane, Canteberra, Melbourne Sydney, Guam, Vladivostock
+11:00	+660		Magadan, Solomon Islands, New Caledonia
+12:00	+720	IDLE - International Date Line East NZST - New Zealand Standard NZT - New Zealand	Wellington, New Zealand; Fiji, Marshall Islands

NTP updates are distributed using GMT time. To properly display the local time in logs and other timestamp information, the switch should be configured with the appropriate offset to GMT based on geographical location.

SNTP Configuration Commands

Table 5.12 describes Simple Network Time Protocol (SNTP) configuration commands. Press the Tab key in the command line interface for more command options.

Table 5.12: SNTP Configuration Commands

Command	Description
configure sntp-client [primary | secondary] server [<ipaddress> | <host_name>]	Configures an NTP server for the switch to obtain time information. Queries are first sent to the primary server. If the primary server does not respond within 1 second, or if it is not synchronized, the switch queries the second server.
configure sntp-client update-interval <seconds>	Configures the interval between polling for time information from SNTP servers. The default setting is `64`.
disable sntp-client	Disables SNTP client functions.
enable sntp-client	Enables SNTP client functions.
show sntp-client	Displays configuration and statistics for the SNTP client.

SNTP Example

In this example, the 480T routing switch queries a specific NTP server and a backup NTP server. An update occurs every 20 minutes. The commands to configure the switch are:

```
configure timezone -480 autodst
configure sntp-client update interval 1200
enable sntp-client
configure sntp-client primary server 10.0.1.1
configure sntp-client secondary server 10.0.1.2
```

6 Configuring Ports

This chapter describes how to configure ports on the Intel® NetStructure™ 480T routing switch and covers these topics:

- Configuring Ports
- Changing Port Speed and Duplex Settings
- Jumbo Frames
- Load Sharing
- Jumbo Frames
- Port-Mirroring
- Enterprise Discovery Protocol

Configuring Ports

By default, all ports are enabled. To enable or disable one or more ports, use this command:

```
[enable | disable] ports <portlist>
```

For example, to disable ports 3, 5, and 12 through 15 on the switch, enter this:

```
disable port 3,5,12-15
```

Even though a port is disabled, the link remains enabled for diagnostic purposes.

Changing Port Speed and Duplex Setting

By default, the switch is configured to use auto-negotiation to determine port speed and duplex setting for each port. You can manually configure the duplex setting and the speed of 100/1000 Mbps ports, and you can manually configure the duplex setting on the GBIC ports

The 480T routing switch fast Ethernet ports can connect to either 100BASE-TX or 1000BASE-T networks. By default, the ports auto-negotiate port speed. You can also configure each port for a particular speed (either 100 Mbps or 1000 Mbps).

The GBIC ports are statically set to 1000 Mbps, and their speed cannot be modified.

To configure port speed and duplex setting, use this command:

```
configure ports <portlist> auto off {speed [100 |
1000]} duplex [half | full]
```

Except for the 10/100 management port, only 100 Mbps and 1000 Mbps speeds are currently supported.

To configure the switch to auto-negotiate, use this command:

```
configure ports <portlist> auto on
```

Flow control is supported only on GBIC ports. It is enabled or disabled as part of auto-negotiation. If auto-negotiation is set to off, flow control is disabled. When auto-negotiation is turned on, flow control is enabled.

Random Early Detection (RED)

Random Early Detection (RED) selectively drops packets when the output interface begins to show signs of congestion. By dropping some packets early rather than waiting until the buffer is full, RED avoids dropping large numbers of packets at once. This minimizes the chance of producing waves of congestion followed by periods when the link is not fully used. Thus, RED allows the transmission line to be used fully at all times. RED statistically drops more packets from large users than small, so traffic sources that generate the most traffic are more likely to be slowed down than traffic sources that generate little traffic.

To turn on RED, use this command:

```
enable red port <portnumber>
```

To configure the probability at which you want random early detection to drop packets, use this command:

```
configure red drop-probability <percent>
```

The percentage range is 0 - 100.

Turning Off Auto-negotiation for a GBIC Port

In certain interoperability situations, it is necessary to turn auto-negotiation off on a GBIC. Even though these ports run only at full duplex and gigabit speeds, the command that turns off auto-negotiation must still include the duplex setting.

This example turns auto-negotiation off for port 4 (a GBIC Mbps Ethernet port):

```
configure ports 4 auto off duplex full
```

Jumbo Frames

Jumbo frames are Ethernet frames that are larger than the allowable maximum size of 1522 bytes, including four bytes used for the cyclic redundancy check (CRC). The switch supports switching and routing of jumbo frames at wire-speed on all ports.

Jumbo frames are used between endstations that support larger frame sizes for more efficient transfers of bulk data. Both endstations (and all the devices in the path) involved in the transfer must be capable of supporting jumbo frames.

Enabling Jumbo Frames

Some network interface cards have a configured maximum MTU size that does not include the additional 4 bytes of CRC. Ensure that the NIC maximum MTU size is at or below the maximum MTU size configured on the switch. Larger frames are dropped at the ingress port.

To enable jumbo frame support, you must configure the MTU size (the largest jumbo frame allowed). To set the MTU size, use this command:

```
configure jumbo-frame size <jumbo_frame_mtu>
```

The jumbo_frame_mtu range is 1523 to 9216. The value describes the maximum size "on the wire," and includes 4 bytes of CRC plus another 4 bytes if 802.1Q tagging is being used.

Next, enable support on the physical ports that will carry jumbo frames, using this command:

```
enable jumbo-frame ports [<portlist> | all]
```

Path MTU Discovery

In path MTU discovery, a source host will assume that the path MTU is the MTU of the first hop, which is known. The host will send all datagrams on that path with the DF (datagram fragmentation) bit set, restricting fragmentation. If any of the datagrams must be fragmented by a switch along the path, that switch will discard the datagrams and return ICMP (Internet Control Message Protocol) Destination Unreachable messages with a code meaning "fragmentation needed and DF set." Upon receipt of such a message (sometimes called a "Datagram Too Big" message), the source host reduces its assumed path MTU and can retransmit.

The path MTU discovery process ends when:

- The host sets the path MTU low enough that its datagrams can be delivered without fragmentation.
- The host does not set the DF bit in the datagram headers.

A host can choose not to set the DF bit because it is willing to have datagrams fragmented. Normally, the host continues to set DF in all datagrams, so that if the route changes and the new PMTU is lower, the host can perform PMTU discovery again.

IP Fragmentation with Jumbo frames

To set the MTU size greater than 1500, all ports in the VLAN must be jumbo-frame enabled.

If an IP packet originates in a local network that allows large packets and that packet traverses a network that limits packets to a smaller size, the packet is fragmented instead of discarded. This is designed for use in conjunction with jumbo frame support.

Frames that are fragmented are not processed at wire-speed within the switch fabric.

Also note that jumbo frame to jumbo frame fragmentation is not supported – only jumbo frame to normal frame fragmentation is currently supported.

To configure VLANs for IP fragmentation:

1. Enable jumbo frames on the incoming port
2. Add the port to a VLAN
3. Assign an IP address to the VLAN
4. Enable IP forwarding on the VLAN
5. Set the MTU size using the following new command:

   ```
   configure ip-mtu <size> vlan <vlan name>
   ```

   The ip-mtu value can be 1500 or 9216, with 1500 the default. If you enter a value other than 1500, the switch will recognize that value as 9216.

IP Fragmentation within a VLAN

The routing switch also supports IP fragmentation within a VLAN. This feature does not require you to configure the MTU size. To use IP fragmentation within a VLAN:

1. Enable jumbo frames on the incoming port
2. Add the port to a VLAN
3. Assign an IP address to the VLAN
4. Enable IP forwarding on the VLAN

If you leave the default MTU (maximum transmission unit) size when you enable jumbo-frame support on a port in the VLAN, you will receive a warning that the VLAN IP-MTU size is not set at maximum jumbo frame size. You can ignore this warning if you want IP fragmentation only within a VLAN. This is for inter-VLAN IP fragmentation only. For intra-VLAN IP fragmentation, all ports in the VLAN must be configured for jumbo frame support.

Load Sharing

Load sharing (also called link aggregation) using 480T routing switches allows you to increase bandwidth and resilience between switches by using a group of ports to carry traffic in parallel between switches. The sharing algorithm allows the routing switch to use multiple ports as a single logical port. For example, VLANs see the load-sharing group as a single logical port. Most load-sharing algorithms also guarantee packet sequencing between clients.

If a port in a round-robin load share group is removed, the traffic that was being transmitted on that link is distributed on only one of the other active load share links in the round-robin group. The traffic is not distributed evenly between the remaining ports.

If a port in a load-sharing group fails, traffic is redistributed to the remaining ports in the load-sharing group. If the failed port becomes active again, traffic is redistributed to include that port.

Load sharing must be enabled on both ends of the link, or a network loop will result.

This feature is supported between 480T routing switches only, but may be compatible with third-party trunking or sharing algorithms. Check with your Intel Customer Service Representative (see "Intel Customer Support" on page 461).

Load Sharing Algorithms

Load-sharing algorithms allow you to select the distribution technique used by the load-sharing group to determine the output port selection. Algorithm selection is not intended for use in predictive traffic engineering.

If you do not explicitly select an algorithm, the port-based scheme is used. However, the address-based algorithm has a more even distribution and is the recommended choice.

You can configure one of three load-sharing algorithms:

- Port-based
- Address-based
- Round-robin

Port-based load sharing algorithms use the ingress port to determine which physical port in the load-sharing group is used to forward traffic out of the switch.

Address-based load-sharing algorithms use addressing information to determine which physical port in the load-sharing group to use for forwarding traffic out of the switch. Addressing information is based on the packet protocol, as follows:

- IP packets—Uses the source and destination MAC and IP addresses, and the TCP port number.
- IPX§ packets—Uses the source and destination MAC address, and IPX network identifiers.
- All other packets—Uses the source and destination MAC address.

Using the round-robin algorithm, packet sequencing between clients is not guaranteed.

Round-robin load-sharing algorithms forward one packet out of each physical port in the load-sharing group using a round-robin scheme, whenever the switch receives a stream of packets.

Configuring Load Sharing

To set up the 480T routing switch to load share among ports, you must create a load-sharing group of ports. The first port in the load-sharing group is configured to be the master logical port. This is the reference port used in configuration commands. You can think of it as if the logical port represents the entire port group.

Do not mix media types such as copper and fiber in a load-sharing configuration.

These rules apply to load sharing:

- A group can contain up to 8 ports.
- The ports in a group do not need to be contiguous.
- Using odd numbered ports (1,3,5,7,9,11) can result in uneven packet distribution across ports.

To define a load-sharing group, you assign a group of ports to a single, logical port number. To enable or disable a load-sharing group, use these commands:

```
enable sharing <port> grouping <portlist>
{algorithm [port-based | address-based | round-
robin]}
```

```
disable sharing <port>
```

Load-Sharing Example

Do not disable a port that is part of a load-sharing group. Disabling the port prevents it from forwarding traffic, but still allows the link to initialize. As a result, a partner switch does not receive a valid indication that the port is not in a forwarding state, and the partner switch will continue to forward packets.

This example defines a load-sharing group that uses ports 9-12, and assigns the first port in the group as the master logical port 9:

```
enable sharing 9 grouping 9-12
```

In this example, logical port 9 represents physical ports 9 through 12.

Always reference the master logical port of the load-sharing group (port 9 in the previous example) when configuring or viewing VLANs. VLANs configured to use other ports in the load-sharing group will have those ports deleted from the VLAN when load sharing becomes enabled.

Verifying the Load Sharing Configuration

The `show ports configuration` command shows whether or not the ports are load sharing and shows the master logical port identity.

Port Commands

Table 6.1 describes the port commands. For further command options, press the Tab key in the command line interface.

Table 6.1: Port Commands

Command	Description
configure jumbo-frame size <jumbo_frame_mtu>	Configures the jumbo frame size. The range is between 1523 and 9216. The default setting is 9216.
configure ports <portlist> auto off {speed [100 | 1000]} duplex [half | full]	Changes the configuration of a group of ports. Specify: • `auto off`—The port will not auto-negotiate the settings. • `speed`—The speed of the port. Except for the 10/100 management port, only 100/1000 speeds are currently supported on the 480T routing switch. • `duplex`—The duplex setting (half-duplex or full-duplex).
configure ports <portlist> auto on	Enables auto-negotiation for the port type; 802.3z for 100/1000 Mbps ports, or 802.3u for the 10/100 management port.
configure ports <portlist> display-string <string>	Configures a user-defined string for a port. The string is displayed in certain `show` commands (for example, `show ports info`). The string can be up to 16 characters.
configure ports [all | mgmt | <portnumber>] qosprofile <qosname>	Configures one or more ports to use a particular QoS profile.
disable jumbo-frame ports [<portlist> | all]	Disables jumbo frame support on a port.
disable learning ports <portlist>	Disables MAC address learning on one or more ports for security purposes. Once disabled, only broadcast traffic, EDP traffic, and packets destined for a permanent MAC address that matches a port number, are forwarded to that port. The default setting is enabled.

Table 6.1: Port Commands (continued)

Command	Description
disable ports <portlist>	Disables a port. Even when disabled, the link is available for diagnostic purposes.
disable sharing <port>	Disables a load-sharing group of ports.
enable jumbo-frame ports [<portlist> | all]	Enables reception and transmission of jumbo frames. A jumbo frame is dropped: • if it is received on a port with jumbo frames disabled, or • if the jumbo frame needs to be forwarded out of a port that has jumbo frames disabled.
enable learning ports <portlist>	Enables MAC address learning on one or more ports. The default setting is enabled.
enable ports [<portlist> | all]	Enables a port.
enable sharing <port> grouping <portlist> {algorithm [port-based | address-based | round-robin]}	Defines a load-sharing group of ports. The ports specified in `<portlist>` are grouped to the master port. Optional load-sharing algorithms include: • `port-based`—Uses the ingress port as criteria for egress port selection. • `address-based`—Uses addressing information as criteria for egress port selection. • `round-robin`—Forwards packets to all egress ports in a `round-robin` fashion. If not specified, port-based load sharing is used.
restart ports {<portlist> | mgmt}	Resets auto-negotiation for one or more ports by resetting the physical link.
show ports {<portlist> | mgmt} collisions	Displays real-time collision statistics.

Table 6.1: Port Commands (continued)

Command	Description
show ports {<portlist> | mgmt} configuration	Displays the port configuration.
show ports {<portlist> | mgmt} info {detail}	Displays detailed system-related information.
show ports {<portlist> | mgmt} packet	Displays a histogram of packet statistics.
show ports {<portlist> | mgmt} qosmonitor	Displays real-time QoS statistics. For more information, refer to "Quality of Service (QoS)" on page 135.
show ports {<portlist> | mgmt} Rxerrors	Displays real-time receive error statistics. For more information on error statistics, refer to "Status Monitoring and Statistics" on page 403.
show ports {<portlist> | mgmt} stats	Displays real-time port statistics.
show ports {<portlist> | mgmt} txerrors	Displays real-time transmit error statistics.
show ports {<portlist> | mgmt} utilization	Displays real-time port utilization information. Use the Spacebar to toggle between packet, byte, and bandwidth utilization information.
unconfigure ports {<portlist> | mgmt } display-string <string>	Clears the user-defined display string from a port.
enable red port <portnumber>	Enables RED on a port.
configure red drop-probability <percent>	Configures the RED drop-probability. The percentage range is 0 - 100.
disable red ports	Disables RED on one or all ports.

Port-Mirroring

Port-mirroring configures the switch to copy all traffic coming in and out of one or more ports to a monitor port on the switch. You can connect the monitor port to a network analyzer or RMON probe for packet analysis.

The switch uses a traffic filter that copies a group of traffic to the monitor port. You can define the traffic filter based on:

- **Physical port**—All data that traverses the port, regardless of VLAN configuration, is copied to the monitor port.
- **VLAN**—All data to and from a particular VLAN, regardless of the physical port configuration, is copied to the monitor port.
- **Virtual port—**All data specific to a VLAN on a specific port is copied to the monitor port.

You can configure up to eight mirroring filters and one monitor port. Once you specify a port as a monitor port, you cannot use it for any other function.

The mirroring port can be tagged or untagged. This allows the mirroring of multiple ports and/or VLANs to a mirror port while preserving the ability of a single protocol analyzer to track and differentiate traffic within a broadcast domain (VLAN) and across broadcast domains (e.g. across VLANS when routing). See "Tagged VLANs" on page 99.

Frames that contain errors are not mirrored.

Mirrored frames that are transmitted from the switch do not contain 802.1Q VLAN tagging information.

Mirroring Combined with Load Sharing

When mirroring ports also involve load-sharing, these limitations apply:

- Mirroring multiple or single VLANs on a specific port is known to cause behavioral problems when used in combination with load sharing. If enabled, load sharing will only make use of the master port and will not fail-over correctly. Deleting the mirror entry will restore normal operation.
- If the master port of a load-shared port group is down, mirroring will not provide traffic for the load-shared port group.

Mirroring IP Multicast Traffic

Due to IGMP snooping, multicast traffic may cease to be seen on a mirror port. If you issue a `restart` command for the mirror port or remove and reinsert the port connection, multicast traffic will resume for the IGMP host time-out period (260 seconds).

Mirroring Bandwidth

Performing mirroring on gigabit ports running at line-rate will reduce the traffic throughput by approximately 30 percent.

Mirroring and Flooding

When a mirrored port is configured, the forwarding database for items being mirrored (e.g. ports or VLANs) are automatically cleared from the forwarding database if the link status on the mirror port changes. This can cause some temporary flooding until the normal learning process completes. Thus, removing or inserting a probe device into the mirror port may appear to cause flooding. However, this is expected behavior.

Mirroring and Download Configuration

When a mirrored port is enabled and configured, a downloaded configuration with mirroring options configured will cause the console to lock up. Manually reset the switch to regain access.

Port-Mirroring Commands

Port-mirroring commands are described in Table 6.2. For further command options, press the Tab key in the command line interface.

Table 6.2: Port-Mirroring Configuration Commands

Command	Description
configure mirroring add [vlan <name> | port <port> | vlan <name> port <port>]	Adds a single mirroring filter definition. You can add up to eight mirroring definitions. You can mirror traffic from a VLAN, a physical port, or a specific VLAN/port combination.

Table 6.2: Port-Mirroring Configuration Commands

Command	Description
configure mirroring delete [vlan <name> | port <port> | vlan <name> port <port>]	Deletes a particular mirroring filter definition, or all mirroring filter definitions.
disable mirroring	Disables port mirroring.
enable mirroring to port <portnumber> [tagged | untagged]	Designates a port as the mirror output port. See "Tagged VLANs" on page 99.
show mirroring	Displays the port-mirroring configuration.

Port-Mirroring Example

This example selects port 3 as the mirror port, and sends all traffic coming into or out of the switch on port 1 to the mirror port:

```
enable mirroring to port 3 untagged
configure mirroring add port 1
```

This next example sends all traffic coming into or out of the switch on port 1 and the VLAN ***default*** to the mirror port (enable mirroring for port 3 first):

```
configure mirroring add port 1 vlan default
```

Enterprise Discovery Protocol

The Enterprise Discovery Protocol (EDP) is used to gather information about neighbor 480T routing switches. EDP is used by the switches to exchange topology information. EDP is also used by the Enterprise Standby Router Protocol (ESRP). Information communicated using EDP includes:

- MAC address (switch ID)
- Software version information
- IP address
- VLAN-IP information
- Port number

EDP Commands

Table 6.3 lists EDP commands. For further command options, press the Tab key in the command line interface.

Table 6.3: EDP Commands

Command	Description
disable edp ports [<portlist> | all]	Disables the EDP on one or more ports.
enable edp ports [<portlist> | all]	Enables generation and processing of EDP messages on one or more ports. The default setting is enabled.

7 Virtual LANs (VLANs)

Setting up Virtual Local Area Networks (VLANs) on the switch eases many time-consuming tasks of network administration while increasing efficiency in network operations.

This chapter describes the concept of VLANs and explains how to implement VLANs on the Intel® NetStructure™ 480T routing switch.

Overview of Virtual LANs

The term VLAN (Virtual Local Area Network) refers to a collection of devices that communicate as if they were on the same physical LAN. Any set of ports (including all ports on the switch) can be considered a VLAN.

You can create up to 3000 VLANs on the Intel® NetStructure™ 480T routing switch.

LAN segments are not restricted by the hardware that physically connects them. The segments are defined by flexible user groups you create with the command-line interface.

Benefits

Implementing VLANs on your networks has several advantages. VLANs:

- Help to control traffic.
- Provide extra security.
- Ease the change and movement of devices.

VLANs Help to Control Traffic

With traditional networks, broadcast traffic can cause congestion, because packets are sent to all network devices, even though the data is not needed by all.

VLANs increase the efficiency of your network because each VLAN can be set up to include only those devices that must communicate with each other.

VLANs Provide Extra Security

Devices within each VLAN can only communicate with member devices in the same VLAN. For a device in VLAN ***Marketing*** to communicate with devices in VLAN ***Sales***, the traffic must cross a routing device specifically configured for that purpose.

VLANs Ease Device Change and Movement

VLANs are not based on physical location. Therefore physical moves of devices do not require manual system updating.

Many network administrators spend much of their time dealing with moves and changes. If users move to a different subnetwork, the addresses of each endstation must be updated manually.

Bi-directional Rate Shaping for Layer 3 Routed VLANs

For more information see "Bi-directional Rate Shaping for Layer 3 Routed VLANs" on page 163.

Bi-directional rate shaping allows you to perform bandwidth management for Layer 2 and Layer 3 traffic flowing both to and from the switch.

You can achieve bi-directional control by defining queue minimum and maximum bandwidth parameters to build true committed information rate capabilities. Also:

- All traffic grouping and bandwidth-management capabilities associated with Quality of Service (QoS) can be used for both directions of traffic.

When switch ports are configured while in Layer 2 mode, MAC block conflicts will not return error messages if Layer 3 mode is later enabled.

- The switch returns error messages on MAC block conflicts when you add rate-shaped ports to VLANs.
- MAC block restrictions do not exist when using the switch as Layer 2 only.

Types of VLANs

You can create VLANs based on these criteria:

- Physical port
- 802.1Q tag
- Ethernet, Logical Link Control Service Advertising Protocol (LLC SAP), or Logical Link Control Subnetwork Access Protocol (LLC/SNAP) Ethernet protocol type
- MAC address
- A combination of these criteria

Port-Based VLANs

In a port-based VLAN, a VLAN name is given to a group of one or more ports on the switch. A port can be a member of only one port-based VLAN.

For example, on the switch in Figure 7.1:

- Ports 1 through 4 are part of VLAN ***Marketing***
- Ports 9 through 12 are part of VLAN ***Sales***
- Ports 5 through 8 and 15 and 16 are in VLAN ***Finance***.

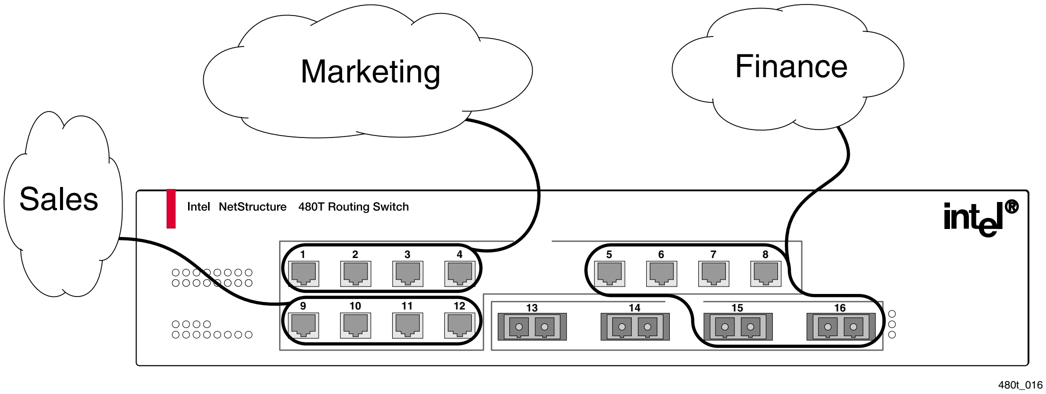

Figure 7.1: Example of a port-based VLAN on the Intel® NetStructure™ 480T routing switch

For the members of the different IP VLANs to communicate, the traffic must be routed by the switch, even if they are physically part of the same port. This means that each VLAN must be configured as a router interface with a unique IP address.

Spanning Switches with Port-Based VLANs

To create a port-based VLAN that spans two switches, you must:

- Assign the port on each switch to the VLAN.
- Connect the two switches using one port on each switch per VLAN.

Figure 7.2 illustrates a single VLAN that spans two 480T routing switches. All ports on both switches belong to VLAN ***Sales***. The two switches are connected using port 13 on System 1, and port 16 on System 2.

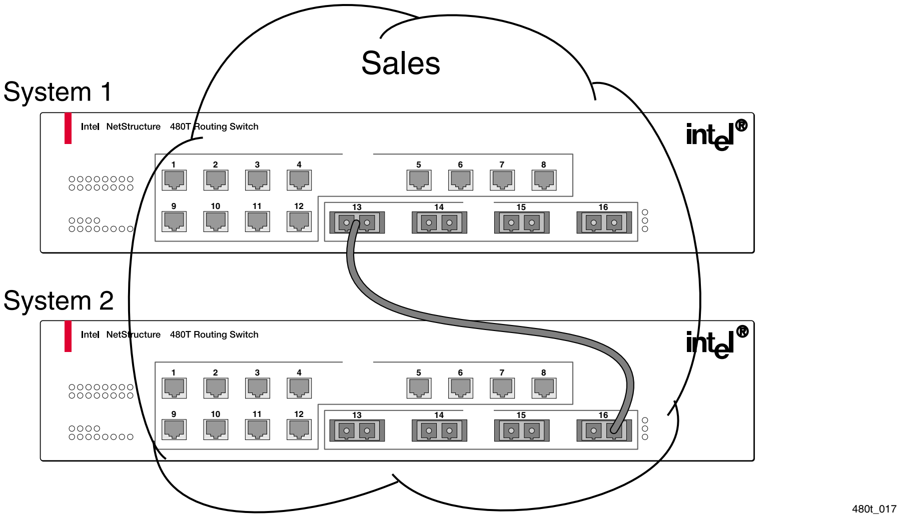

Figure 7.2: Single port-based VLAN spanning two switches

To create multiple VLANs that span two switches in a port-based VLAN, a port on System 1 must be connected to a port on System 2 for each spanned VLAN and each of these ports must also be a member of the corresponding VLANs.

Figure 7.3 illustrates two VLANs spanning two switches:

- On System 1, ports 9 through 12 are part of VLAN ***Accounting*** and ports 13 through 16 are part of VLAN ***Engineering***.

- On System 2, ports 1 through 4 are part of VLAN ***Accounting*** and ports 5 through 8, 15, and 16 are part of VLAN ***Engineering***.

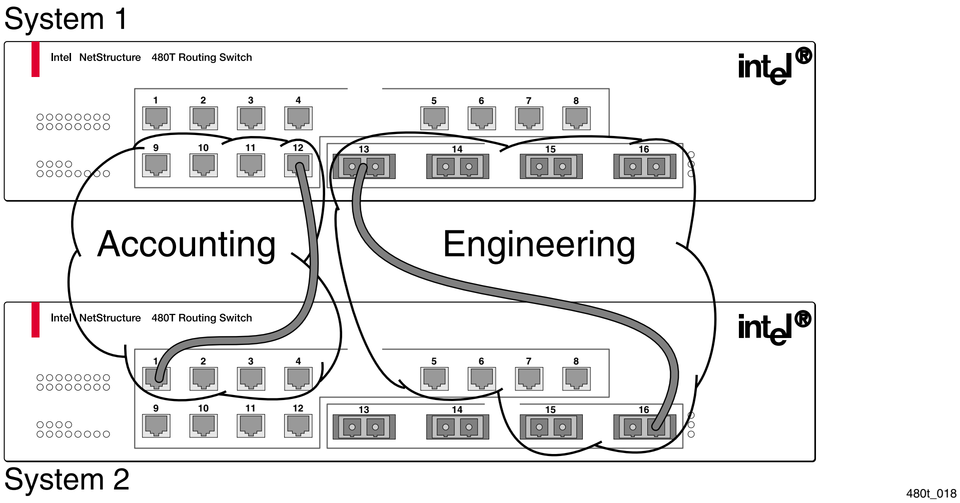

Figure 7.3: Two port-based VLANs spanning two switches

- VLAN ***Accounting*** spans System 1 and System 2 by way of a connection between System 1, port 12 and System 2, port 1.
- VLAN ***Engineering*** spans System 1 and System 2 by way of a connection between System 1, port 13, and System 2, port 16.

Using this configuration, you can create multiple VLANs that span multiple switches, in a daisy-chained fashion. To function properly:

- Each switch must have a dedicated port for each VLAN.
- Each dedicated port must be connected to a port that is a member of its VLAN on the next switch.

Tagged VLANs

Tagging is a process that inserts a marker (called a tag) into the Ethernet frame. The tag includes the identification number of a specific VLAN, called the VLANid.

Using 802.1Q tagged packets may create packets slightly bigger than the current IEEE 802.3/Ethernet maximum of 1,518 bytes. This may affect packet error counters in other devices, and may also lead to connectivity problems if non-802.1Q bridges or routers are placed in the path.

Uses of Tagged VLANs

Tagging is most commonly used to create VLANs that span switches. The switch-to-switch connections are typically called trunks. Using tags, multiple VLANs can span multiple switches using one or more trunks. In a port-based VLAN, each VLAN requires its own pair of trunk ports, as shown in Figure 7.3. Using tags, multiple VLANs can span two switches with a single trunk.

Another benefit of tagged VLANs is that a port can be a member of multiple VLANs. This is particularly useful if you have a device (such as a server) that must belong to multiple VLANs.

A single port can be a member of only one port-based VLAN. If you want the port to be a member of more than one VLAN, it must use tagging. All additional VLAN membership for the port must be accompanied by tags. Along with configuring the VLAN tag for the port, the server must have a network interface card (NIC) that supports 802.1Q tagging.

Assigning a VLAN Tag

Each VLAN may be assigned an 802.1Q VLAN tag. As ports are added to a VLAN with an 802.1Q tag defined, you decide whether each port will use tagging for that VLAN. The default mode of the switch is to have all ports assigned to the VLAN named ***default*** with an 802.1Q VLAN tag (VLANid) of 1 assigned.

Any packets arriving tagged with a VLANid that is not configured on a port is discarded.

Not all ports in the VLAN must be tagged. As traffic from a port is forwarded out of the switch, the switch determines (in real time) if each destination port should use tagged or untagged packet formats for that VLAN. The switch adds and deletes tags, as required, by the port configuration for that VLAN.

Figure 7.4 illustrates the physical view of a network that uses tagged and untagged traffic.

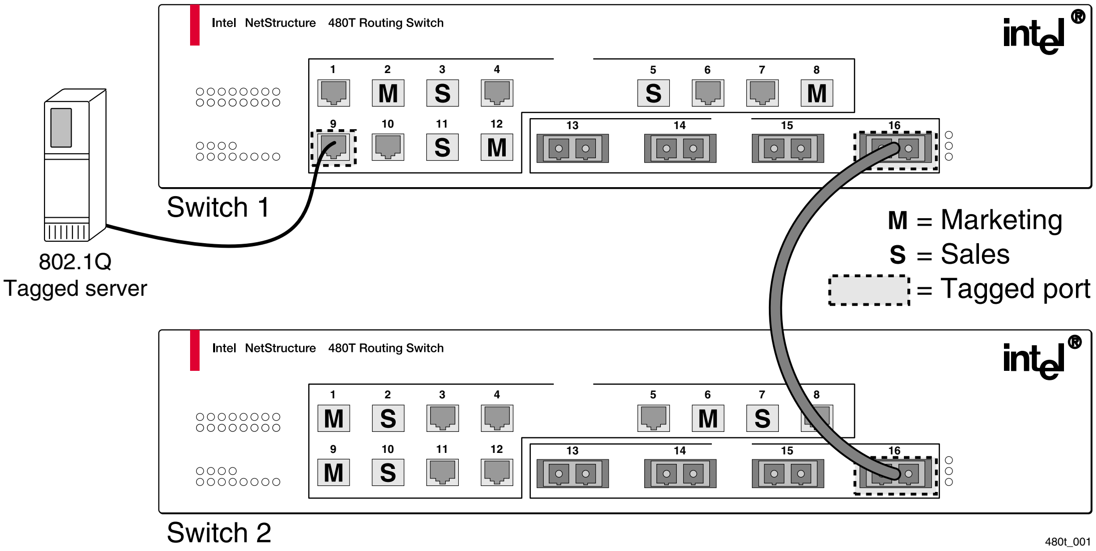

Figure 7.4: Physical diagram of tagged and untagged traffic

Figure 7.5 shows a logical diagram of the same network.

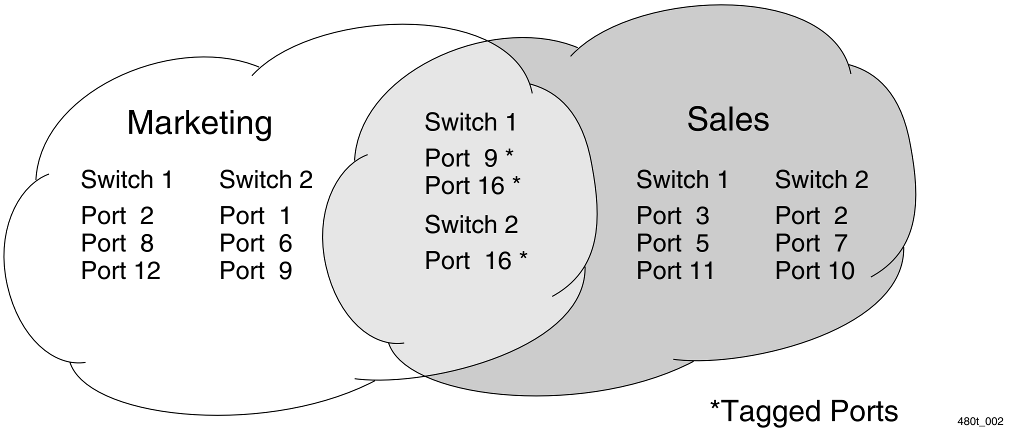

Figure 7.5: Logical diagram of tagged and untagged traffic

In Figure 7.4 and Figure 7.5:

- The trunk port on each switch carries traffic for both VLAN ***Marketing*** and VLAN ***Sales***.

- The trunk port on each switch is tagged.
- The server connected to port 9 on System 1 has a NIC that supports 802.1Q tagging.
- The server connected to port 9 on System 1 is a member of both VLAN ***Marketing*** and VLAN ***Sales***.
- All other stations use untagged traffic.

As data passes out of the switch, the switch determines if the destination port requires the frames to be tagged or untagged. All traffic coming from and going to the server is tagged. Traffic coming from and going to the trunk ports is tagged. The traffic that comes from and goes to the other stations on this network is not tagged.

Mixing Port-Based and Tagged VLANs

You can configure the switch using a combination of port-based and tagged VLANs. Each port can be a member of multiple VLANs, with the stipulation that only one of its VLANs uses untagged traffic. In other words, a port can simultaneously be a member of one port-based VLAN and multiple tag-based VLANs.

VLAN classification treats packets arriving on a port with an 802.1Q tag containing a VLANid of zero as untagged.

Protocol-Based VLANs

Protocol-based VLANs allow you to define a packet filter as the matching criteria to determine if a packet belongs to a particular VLAN.

Protocol-based VLANs are most often used in situations where network segments include hosts running multiple protocols. For example, in Figure 7.6, the hosts are running both the IP and NetBIOS§ protocols:

- The IP traffic is divided into two IP subnets, 192.207.35.0 and 192.207.36.0.
- The subnets are internally routed by the switch.
- The subnets are assigned different VLAN names, ***Finance*** and ***Personnel***, respectively.

- The remainder of the traffic belongs to the VLAN named ***MyCompany***.
- All ports are members of the VLAN ***MyCompany***.

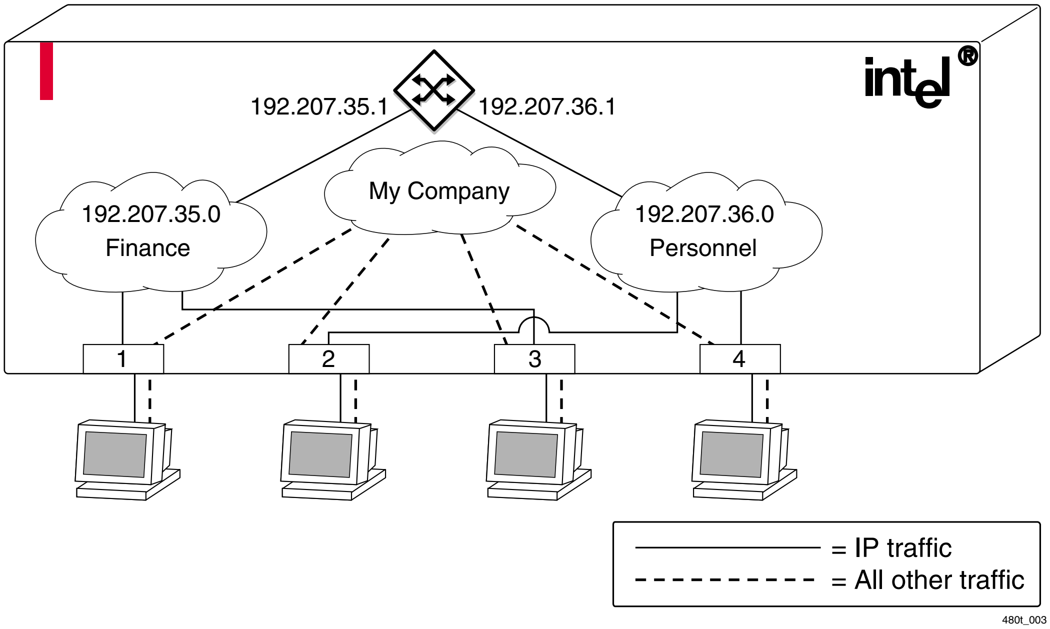

Figure 7.6: Protocol-based VLANs

Predefined Protocol Filters

These protocol filters are predefined on the switch:

- IP
- IPX§
- NetBIOS
- DECnet§
- IPX_8022
- IPX_SNAP
- AppleTalk§

Defining Protocol Filters

For more information on SNAP for Ethernet protocol types, see TR 11802-5:1997 (ISO/IEC) [ANSI/IEEE std. 802.1H, 1997 Edition]. For more information on standards see "Technical Specifications and Supported Limits" on page 431.

If necessary, you can define a customized protocol filter, based on Ethertype, Logical Link Control (LLC), and/or Subnetwork Access Protocol (SNAP). Up to six protocols may be part of a protocol filter.

To define a protocol filter:

1. Create a protocol using this command:

   ```
   create protocol <protocol_name>
   ```

   For example:

   ```
   create protocol fred
   ```

   The protocol name can have a maximum of 32 characters.

2. Configure the protocol using this command:

   ```
   configure protocol <protocol_name> add
   <protocol_type> <hex_value>
   ```

Supported protocol types include:

- **etype**—Ethertype

  The values for **etype** are four-digit hexadecimal numbers taken from a list maintained by the IEEE. You can find this list at this URL:

  http://standards.ieee.org/regauth/ethertype/index.html

- **llc**—LLC Service Advertising Protocol (SAP)

  The values for **LLC** are four-digit hexadecimal numbers that are created by concatenating a two-digit LLC Destination SAP (DSAP) and a two-digit LLC Source SAP (SSAP).

- **snap**—Ethertype inside an IEEE SNAP packet encapsulation.

  The values for **snap** are the same as the values for **etype**, described previously.

  For example:

  ```
  configure protocol fred add llc feff
  configure protocol fred add snap 9999
  ```

You can define a maximum of fifteen protocol filters, each containing a maximum of six protocols. All fifteen protocol filters can be active and configured for use.

Deleting a Protocol Filter

If a protocol filter is deleted from a VLAN, the VLAN is assigned a protocol filter of `none`. You can continue to configure the VLAN. However, no traffic is forwarded to the VLAN until a protocol is assigned to it.

Precedence of Tagged Packets Over Protocol Filters

If a VLAN is configured to accept tagged packets on a particular port, incoming packets that match the tag configuration take precedence over any protocol filters associated with the VLAN.

VLAN Names

Each VLAN is given a name that can be up to 32 alphanumeric characters.

VLAN names normally begin with an alphabetical letter. Use quotation marks to enclose a VLAN name that does not begin with an alphabetical character, or that includes a space, comma, or other special character. For example:-

Table 7.1:

Correct	Incorrect
vlanmarketing2 “2ndvlangroup” “vlan marketing”	2ndvlangroup vlan marketing

Use VLAN names consistently across your entire network.

VLAN names are locally significant. That is, VLAN names used on one switch are only meaningful to that switch. If another switch is connected to it, the VLAN names have no significance to the other switch.

Default VLAN

The switch ships with one default VLAN that has these properties:

- The VLAN name is ***default.***
- It includes all the ports on a new or initialized switch.

The default VLAN is untagged on all ports. It has an internal VLANid of 1.

Renaming a VLAN

To rename a VLAN, use this command:

```
configure vlan <old_name> name <new_name>
```

These rules apply to renaming VLANs:

- Once you change the default VLAN name, it cannot be changed back to ***default***.
- You cannot create a new VLAN named ***default***.
- You cannot change the VLAN name ***MacVlanDiscover***. Although the switch accepts a name change, the original name is recreated after the switch is rebooted.

Configuring VLANs on the Switch

This section describes the commands associated with setting up VLANs on the switch.

To configure a VLAN:

1. Create and name the VLAN.
2. Assign an IP address and mask (if applicable) to the VLAN, if needed.
3. Assign a VLANid, if any ports in this VLAN uses a tag.
4. Assign one or more ports to the VLAN.

Each IP address and mask assigned to a VLAN must represent a unique IP subnet. You cannot configure the same IP subnet on different VLANs.

As you add each port to the VLAN, decide if the port uses an 802.1Q tag.

Table 7.2 describes the commands used to configure a VLAN. For a complete list of command options, press the Tab key in the command line interface.

Table 7.2: VLAN Configuration Commands

Command	Description
configure dot1q ethertype <ethertype>	Configures an IEEE 802.1Q Ethertype. Use this command only if you have another switch that supports 802.1Q, but uses an Ethertype value other than 8100. You must reboot the switch for this command to take effect.
configure protocol <protocol_name> [add | delete] <protocol_type> <hex_value> {<protocol_type> <hex_value>} ...	Configures a protocol filter. Supported `<protocol_type>` values include: • `etype` • `llc` • `snap` The variable `<hex_value>` is a hexadecimal number between 0 and FFFF that represents either the Ethernet protocol type, the Destination Service Access Point/Session Service Access Point (DSAP/SSAP) combination (for Link Level Control (LLC)), or the Subnetwork Network Access Protocol (SNAP)-encoded Ethernet protocol type (for SNAP).
configure vlan <name> add port [<portlist> | all] {tagged | untagged} {nobroadcast}	Adds one or more ports to a VLAN. You can specify tagged or untagged ports. Specify `nobroadcast` to prevent the forwarding of broadcast, multicast, and unknown unicast traffic. By default, ports are `untagged`.
configure vlan <name> delete port [<portlist> | all]	Deletes one or more ports from a VLAN.
configure vlan <name> ipaddress <ipaddress> {<mask>}	Assigns an IP address and an optional mask to the VLAN.

Table 7.2: VLAN Configuration Commands (continued)

Command	Description
configure vlan <name> protocol [<protocol_name> | any]	Configures a protocol-based VLAN. If the keyword `any` is specified, it becomes the default VLAN. All packets that cannot be forwarded to other protocol-based VLANs are assigned to the default VLAN of that port.
configure vlan <name> qosprofile [<qosprofile> | none]	Configures a VLAN to use a particular QoS profile. Dynamic forwarding database entries (FDB) associated with the VLAN are flushed once the change is committed.
configure vlan <name> tag <vlanid>	Assigns a numerical VLANid. The valid range is from 1 to 4095.
configure vlan <old-name> name <new-name>	Changes the name of a configured VLAN.
create protocol <protocol_name>	Creates a user-defined protocol.
create vlan <name>	Creates a named VLAN.
delete protocol <protocol>	Removes a protocol.
delete vlan <name>	Removes a VLAN.
unconfigure vlan <name> ipaddress	Resets the VLAN IP address.
unconfigure vlan <name> xnetid	Resets the VLAN xnetid.

VLAN Configuration Examples

Example 1

This example creates a port-based VLAN named ***accounting***, assigns the IP address 132.15.121.1, and assigns ports 1, 2, 3 and 6 to it:

```
create vlan accounting
```

```
configure accounting ipaddress 132.15.121.1
configure default delete port 1-3,6
```

```
configure accounting add port 1-3,6
```

Because VLAN names are unique, you do not need to enter the keyword `vlan` after you have created the unique VLAN name. You can use the VLAN name alone.

Example 2

This example creates a tag-based VLAN named ***video***. It assigns the VLANid 1000. Ports 4 through 8 are added as tagged ports to the VLAN.

```
create vlan video
configure video tag 1000
configure video add port 4-8 tagged
```

Example 3

This example creates a VLAN named ***sales***, with the VLANid 120. The VLAN uses both tagged and untagged ports. Ports 1 through 3 are tagged, and ports 4 and 7 are untagged. When not explicitly specified, ports are added as untagged.

```
create vlan sales
configure sales tag 120
configure sales add port 1-3 tagged
configure sales add port 4,7
```

Example 4

This example creates a protocol-based VLAN named ***ipsales***. Ports 6 through 8 are assigned to the VLAN.

```
create vlan ipsales
configure ipsales protocol ip
configure ipsales add port 6-8
```

Example 5

This example defines a protocol filter, ***myprotocol*** and applies it to the VLAN named ***myvlan***. This is a command syntax example only, and has no real-world application.

```
create protocol myprotocol
configure protocol myprotocol add etype 0xf0f0
configure protocol myprotocol add etype 0xffff
create vlan myvlan
configure myvlan protocol myprotocol
```

Displaying VLAN Settings

To display VLAN settings, use this command:

```
show vlan {<name>}
```

The `show` command displays summary information about each VLAN, and includes:

- Name
- VLANid
- How the VLAN was created
- IP address
- IPX address (if configured)
- STPD information
- Protocol information
- QoS profile information
- Ports assigned
- Tagged/untagged status for each port
- How the ports were added to the VLAN

If you want to show all VLANs, type:

```
show vlan detail
```

To display protocol information, use this command:

```
show protocol {<protocol>}
```

This `show` command displays protocol information, including:

- Protocol name
- List of protocol fields

VLAN Statistics

You can collect statistics on a per VLAN basis. Available statistics include:

- Receive and Transmit Unicast
- Receive and Transmit Multicast
- Receive and Transmit Broadcast
- Receive and Transmit Byte Count.

To display VLAN statistics use the command:

```
show vlan stats vlan <vlan_name> <vlan_name>
```

You can use multiple VLAN names in this syntax for multiple VLAN displays.

Deleting VLANs

To delete a VLAN, or to return VLAN settings to their defaults, use the commands listed in Table 7.3. For a complete list of command options, press the Tab key in the command line interface.

Table 7.3: VLAN Delete and Reset Commands

Command	Description
delete protocol <protocol>	Removes a protocol.
delete vlan <name>	Removes a VLAN.
unconfigure vlan <name> ipaddress	Resets the IP address of the VLAN.

VLAN Tunneling (vMANs)

Tunneling technology (also called encapsulating) allows you to send data from one network through another network's connections. It does this by encapsulating a network protocol within data packets carried by the second network.

You can tunnel any number of 802.1Q VLANs into a single VLAN that can be switched through the 480T routing switch Ethernet infrastructure.

Each tunnel is completely isolated from other tunnels or VLANs. This feature is useful in building transparent private networks (also called virtual metropolitan area networks or vMANs) that need point-to-point or point-to-multipoint connectivity across an Ethernet infrastructure.

The VLAN tagging methods used within the vMAN tunnel are transparent to the tunnel. The tagging numbers and methods used by the customer are transparent to the metropolitan area network (MAN) provider.

To configure a vMAN tunnel:

1. Modify the 802.1Q Ethertype used by the switch to recognize tagged frames.
2. Configure the switch to accept larger MTU (maximum transmission unit) size frames (jumbo frames).
3. Create tunnels by creating VLANs and configuring member ports as tagged on switch-to-switch ports and untagged on the tunnel's ingress/egress ports.

The figure shows a vMAN configuration with two tunnels that have ingress/egress ports on each 480T routing switch.

Figure 7.7: vMAN Configuration

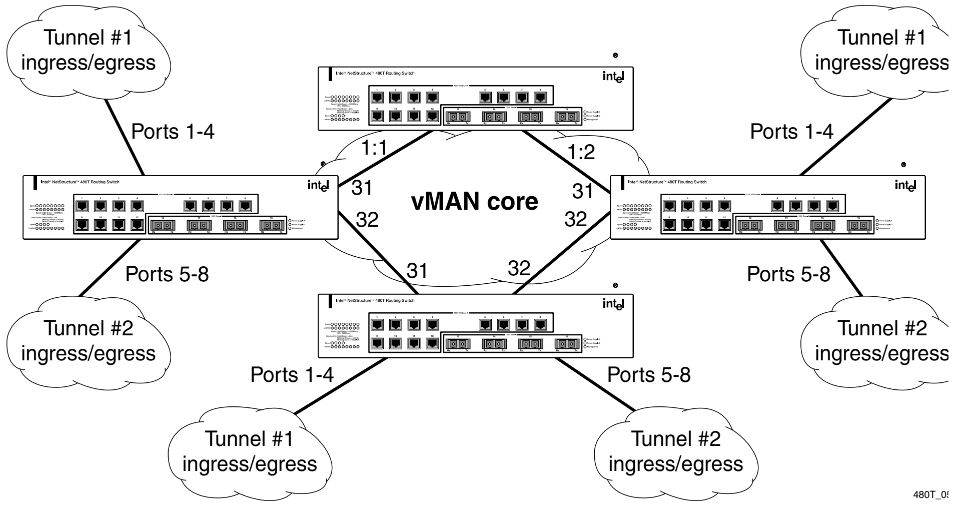

The switches are configured as follows:

```
configure dot1q ethertype 9100

enable jumbo-frame ports 1,2

configure jumbo-frame size 1530

create vlan Tunnel1

configure vlan Tunnel1 tag 50

configure vlan Tunnel1 add port 1-4 untag

configure vlan Tunnel1 add port 1,2 tagged

create vlan Tunnel2

configure vlan Tunnel2 tag 60

configure vlan Tunnel2 add port 5-8 untag

create vlan Tunnel2 add port 1,2 tagged
```

Specific to this configuration, a Layer 1 or Layer 2 redundancy method would also be employed, such as Spanning Tree or other protocol available on the switch.

MAC-Based VLANs

MAC-based VLANs allow physical ports to be mapped to a VLAN based on the source MAC address learned in the forwarding database (FDB). This feature allows you to designate a set of ports that have their VLAN membership dynamically determined by the MAC address of the endstation that plugs into the physical port.

You can configure the source MAC address-to-VLAN mapping either offline or dynamically on the switch.

For example, you can use this application for a roaming user who wishes to connect to a network from a conference room. In each conference room, the user plugs into one of the designated ports on the switch and is mapped to the appropriate VLAN. Connectivity is maintained to the network with all of the benefits of the configured VLAN in terms of QoS, routing, and protocol support.

MAC-Based VLAN Guidelines

When using MAC-to-VLAN mapping, consider these guidelines:

- A port can only accept connections from an endstation/host and should not be connected to a Layer-2 repeater device.
- Connecting to a Layer-2 repeater device can prevent certain addresses from mapping to their respective VLAN if they are not correctly configured in the MAC-VLAN configuration database.

  If a repeater device is connected to a MAC-based VLAN port, and the configured MAC-to-VLAN mapped station enters on the repeater, any endstation that is attached to the repeater can be mapped to that VLAN while the configured endstation is active in that VLAN. After removing the configured MAC-to-VLAN endstation, all other endstations lose connectivity.
- Groups are used as a security measure to allow a MAC address to enter into a VLAN only when the group mapping matches the port mapping.

This example show MAC 00:00:00:00:00:aa is only allowed to enter into the VLAN on ports 10 and 11 because of membership in group 100:

```
* switch480T:50 # show mac
Port                         Vlan            Group          State
10                           MacVlanDiscover 100            Discover
11                           MacVlanDiscover 100            Discover
12                           MacVlanDiscover any            Discover
13                           MacVlanDiscover any            Discover
14                           MacVlanDiscover any            Discover
Total Entries in Database:2
Mac                          Vlan            Group
00:00:00:00:00:aa            sales           100
00:00:00:00:00:01            sales           any
2 matching entries
```

- The group ***any*** is equivalent to the group 0 (zero). Ports that are configured as any allow any MAC address to be assigned to a VLAN, regardless of group association.
- You can download partial configurations of the MAC-to-VLAN database to the switch using the timed download configuration feature. See "Timed Configuration Download, MAC-Based VLANs" on page 117 for more information.

MAC-Based VLAN Limitations

The limitations of MAC-based VLANs are:

You can download up to 7,000 MAC addresses to the switch when using MAC-based VLANs.

- Ports participating in MAC VLANs must first be removed from any static VLANs.
- The MAC-to-VLAN mapping can only be associated with VLANs that exist on the switch.
- A MAC address cannot be configured to associate with more than one VLAN. If this is attempted, the MAC address is associated with the most recent VLAN entry in the MAC-to-VLAN database.
- The feature is intended to support one client per physical port. After a client MAC address has successfully registered, the

VLAN association remains until the port connection is dropped or the FDB entry ages out.

MAC-Based VLAN Commands

Table 7.4 describes MAC-based VLAN commands. For a complete list of command options, press the Tab key in the command line interface.

Table 7.4: MAC-Based VLAN Commands

Command	Description
configure mac-vlan add mac-address [any | <mac_address>] mac-group [any | <group_number>] vlan <name>	Adds a MAC address to a MAC-based VLAN.
configure mac-vlan delete [mac-address <mac_address> | all]	Removes one or all MAC addresses from a MAC-based VLAN.
disable mac-vlan port <portlist>	Disables a port from using the MAC-based VLAN algorithm.
enable mac-vlan mac-group [any | <group_number>] port <portlist>	Enables a port to use the MAC-based VLAN algorithm.
show mac-vlan {configuration | database}	Displays the MAC-based VLAN configuration and MAC address database content.

MAC-Based VLAN Example

In the following example, three VLANs are created, named ***engineering***, ***marketing***, and ***sales***:

- A single MAC address is associated with each VLAN.
- The MAC address 00:00:00:00:00:02 has a group number of ***any*** or 0 (zero) associated with it, allowing it to be inserted into any port that is in MacVlanDiscover mode (ports 1-4 in this case).

- The MAC address 00:00:00:00:00:01 has a group number of 10 associated with it, and can only be assigned to a VLAN if inserted into ports 5 or 6.
- The MAC address 00:00:00:00:00:03 has a group number of 200 associated with it and can only be inserted into ports 9 through 12.

To create the VLANs use these commands:

```
enable mac-vlan mac-group any ports 1-4
enable mac-vlan mac-group 10 ports 5-6
enable mac-vlan mac-group 200 ports 9-12

configure mac-vlan add mac-address
00:00:00:00:00:01 mac-group 10 engineering

configure mac-vlan add mac-address
00:00:00:00:00:02 mac-group any marketing

configure mac-vlan add mac-address
00:00:00:00:00:03 mac-group 200 sales
```

Timed Configuration Download, MAC-Based VLANs

To allow centralized control of MAC-based VLANs over multiple switches, a timed TFTP configuration download allows you to download incremental configuration files from a primary or secondary server at specified time intervals. The timed downloads are configurable in 24-hour intervals. When a switch reboots, the configuration is automatically downloaded according to primary and secondary server settings.

To configure the primary and/or secondary server and file name, use this command:

```
configure download server [primary | secondary]
<host_name> | <ip_address> <filename>
```

To enable timed interval downloads, use this command:

```
download configuration every <hour> <min>
```

To display timed download information, use this command:

```
show switch
```

Example

For MAC-based VLANs, the downloaded file is an ASCII file that consists of CLI commands used to configure the most recent MAC-to-VLAN database.

This feature is different from the normal download configuration command in that it allows incremental configuration without automatically rebooting.

This example shows an incremental configuration file for MAC-based VLAN information that updates the database and saves changes:

```
configure mac-vlan add mac-address
00:00:00:00:00:01 mac-group any engineering

configure mac-vlan add mac-address
00:00:00:00:ab:02 mac-group any engineering

configure mac-vlan add mac-address
00:00:00:00:cd:04 mac-group any sales

configure mac-vlan add mac-address
00:00:00:00:ab:50 mac-group any sales

configure mac-vlan add mac-address
00:00:00:00:cd:60 mac-group any sales

save
```

8 Forwarding Database (FDB)

This chapter describes the contents of the forwarding database (FDB), how the FDB works, and how to configure the FDB.

Overview of the FDB

The Intel® NetStructure™ 480T routing switch maintains a database of all media access control (MAC) addresses received on all of its ports. It uses the information in this database to decide whether a frame should be forwarded or filtered.

IP FDB Performance

The IP FDB handling is enhanced so that only relevant IP FDB entries are flushed when entries are modified in the system routing table.

As a result, you will see a significant performance improvement in situations where there are frequent route changes. Performance is improved because route changes do not affect traffic that is not relevant to the route change.

The 480T routing switch supports 256K entries in the forwarding database. These can be Layer 2 or Layer 3 addresses.

Up to 256 static MAC entries are supported.

You can download up to 7,000 MAC addresses to the switch when using MAC-based VLANs. You can create up to 3,000 VLANs on the switch.

FDB Contents

Each FDB entry consists of:

- The MAC address of the device
- An identifier for the port on which it was received
- An identifier for the VLAN to which the device belongs.

Frames destined for devices that are not in the FDB are flooded to all members of the VLAN.

FDB Entry Types

There are four types of entries in the FDB:

- Dynamic entries
- Non-aging entries
- Permanent entries
- Blackhole entries

Dynamic Entries

Initially, all entries in the database are dynamic.

Entries in the database are removed (aged-out) if, after a period of time (aging time), the device has not transmitted. This prevents the database from filling with obsolete entries by deleting the entry when a device is removed from the network.

Dynamic entries are deleted from the database if the 480T routing switch is reset or a power off/on cycle occurs.

For information about setting the aging time, refer to "Configuring FDB Entries" on page 122.

Non-aging Entries

If the aging time is set to zero, all aging entries in the database are defined as static, non-aging entries. This means that they do not age, but they are still deleted if the switch is reset.

Permanent Entries

All entries entered through the command line interface are stored as permanent.Only entries designated as Permanent are retained in the database if the switch is reset or a power off/on cycle occurs. A permanent entry can either be a unicast or multicast MAC address.

The switch can support up to 256 permanent MAC entries in the forwarding database.

Once created, permanent entries cannot be updated. For example, the permanent entry store is not updated when any of the following take place:

- A VLAN is deleted.
- A VLANid is changed.
- A port mode is changed (tagged/untagged).
- A port is deleted from a VLAN.
- A port is disabled.
- A port enters a blocking state.
- A port QoS setting is changed.
- A port goes down (link down).

Blackhole Entries

A blackhole entry configures the 480T routing switch to discard packets with a specified MAC destination address.

Blackhole entries are useful as a security measure or in special circumstances where a specific destination address must be discarded.

Blackhole entries are treated like permanent entries in the event of a switch reset or power off/on cycle.

Blackhole entries are never aged out of the database.

How FDB Entries Get Added

Add entries to the FDB in two ways:

- The switch can learn entries by updating its FDB with the source MAC address from a packet, the VLAN, and the port identifier where the source packet was received.

- You can enter and update entries using a MIB browser, an SNMP Network Manager, or the command line interface (CLI).

Associating a QoS Profile with an FDB Entry

The switch applies the QoS profile as soon as the FDB entry is learned.

You can associate a QoS profile with a MAC address (and VLAN) of a device that is dynamically learned. The FDB treats the entry like a dynamic entry (it is learned, it can be aged out of the database, and so on).

Configuring FDB Entries

To configure entries in the FDB, use the commands listed in Table 8.1. For further command options, press the Tab key in the command line interface.

Table 8.1: FDB Configuration Commands

Command	Description
create fdbentry <mac_address> vlan <name> [blackhole | ports [<portlist> | all] | dynamic] {qosprofile <qosprofile>}	Creates an FDB entry. Specify: • **mac_address**—Device MAC address, using bytes separated by colons. • **name**—VLAN associated with MAC address. • **blackhole**–Configures the MAC address as a blackhole entry. • **portlist**—Port numbers associated with MAC address. • **dynamic**—Specifies that the entry is learned dynamically. Used to associate a QoS profile with a dynamically learned entry. • **qosprofile**—QoS profile associated with MAC address. If more than one port number is associated with a permanent MAC entry, packets are multicast to the multiple destinations.

Table 8.1: FDB Configuration Commands (continued)

Command	Description
configure fdb agingtime <number>	Configures the FDB aging time (in seconds). The range is 15 through 1,000,000. The default value is `300`. A value of 0 indicates that the entry is never aged out.
disable learning port <portlist>	Disables MAC-address learning on one or more ports for security purposes. Once disabled, only broadcast traffic, EDP traffic, and packets destined for a permanent MAC address that matches the port number, are forwarded to the port. The default setting is enabled.
enable learning port <portlist>	Enables MAC-address learning on one or more ports.

FDB Configuration Examples

This example adds a permanent entry to the FDB:

```
create fdbentry 00:A0:C9:12:34:56 vlan marketing
port 4
```

The permanent entry has these characteristics:

- MAC address is 00:A0:C9:12:34:56.
- VLAN name is ***marketing***.
- Port number for this device is 4.

This example associates the QoS profile ***qp2*** with a dynamic entry that is learned by the FDB:

If you assign a MAC address to a port that is not part of a VLAN, that port will be configured as a black hole.

```
create fdbentry 00:A0:C9:12:34:56 vlan net34
dynamic qosprofile qp2
```

This entry has these characteristics:

- MAC address is 00:A0:C9:12:34:56.
- VLAN name is ***net34***.
- The entry is learned dynamically.
- QoS profile ***qp2*** is applied when the entry is learned.

Displaying FDB Entries

To display FDB entries, use the command:

```
Show fdb {<mac_address> | vlan <name> | ports
<portlist> | permanent}
```

where the following is true:

- `mac_address`—Displays the entry for a particular MAC address.
- `vlan <name>`—Displays the entries for a VLAN.
- `portlist`—Displays the entries for a port.
- `permanent`—Displays all permanent entries.

With no options, the command displays all FDB entries.

Removing FDB Entries

You can remove one or more specific entries from the FDB, or you can clear the entire FDB of all entries by using the commands listed in Table 8.2. For further command options, press the Tab key in the command line interface.

Table 8.2: Removing FDB Entry Commands

Command	Description
clear fdb {<mac_address> l vlan <name> l ports <portlist>}	Clears dynamic FDB entries that match the filter. When no options are specified, the command clears all FDB entries.
delete fdbentry <mac_address> vlan <name>	Deletes a permanent FDB entry.

9 Spanning Tree Protocol (STP)

Using the Spanning Tree Protocol (STP) functionality of the Intel® NetStructure™ 480T routing switch makes your network more fault tolerant.

STP is a part of the 802.1D bridge specification defined by the IEEE (Institute of Electrical and Electronics Engineers), a standard-setting body. To explain STP in terms used by the 802.1D specification, the switch is referred to as a bridge.

Overview of Spanning Tree Protocol

STP is a bridge-based mechanism for providing fault tolerance on networks. STP allows you to implement parallel paths for network traffic, and ensure that the redundant paths are:

- Disabled when the main paths are operational.
- Enabled when the main path fails.

Spanning Tree Domains

You can partition the switch into multiple virtual bridges. Each virtual bridge can run an independent Spanning Tree instance. Each Spanning Tree instance is called a Spanning Tree Protocol Domain (STPD). Each

STPD has its own Root Bridge and active path. After the STPD is created, you can assign one or more VLANs to it.

A port can belong to only one STPD. If a port is a member of multiple VLANs, then all those VLANs must belong to the same STPD.

Remember these key points when configuring VLANs and STP:

- Each VLAN forms an independent broadcast domain.
- STP blocks paths to create a loop-free environment.
- When STP blocks a path, no data can be transmitted or received on the blocked port.
- Within any given STPD, all member VLANs use the same spanning tree.

Be sure that multiple STPD instances within a single switch do not see each other in the same broadcast domain. This could happen if, for example, another external bridge is used to connect VLANs belonging to separate STPDs.

If you delete an STPD, the VLANs that were members of that STPD are also deleted. You must remove all VLANs associated with the STP before deleting the STPD.

In order for the switch to pass bridge protocol data units (BPDUs), with spanning tree disabled, a protocol filter of any must be associated with a VLAN on a port connected to the spanning tree domain. Otherwise when STP is off, BPDUs will not be flooded to adjacent bridges.

STP Configurations

When you assign VLANs to an STPD, pay careful attention to the STP configuration and its effect on the forwarding of VLAN traffic.

Figure 9.1 illustrates a network that uses VLAN tagging for trunk connections.

Five VLANs have been defined:

- ***Sales*** is defined on Switch A, Switch B, and Switch M.
- ***Personnel*** is defined on Switch A, Switch B, and Switch M.
- ***Manufacturing*** is defined on Switch Y, Switch Z, and Switch M.
- ***Engineering*** is defined on Switch Y, Switch Z, and Switch M.

- ***Marketing*** is defined on all switches (Switch A, Switch B, Switch Y, Switch Z, and Switch M).

Two STPDs are defined:

- STPD1 contains VLANs ***Sales*** and ***Personnel.***
- STPD2 contains VLANs ***Manufacturing*** and ***Engineering.***

The VLAN ***Marketing*** is a member of the default STPD, but not assigned to either STPD1 or STPD2.

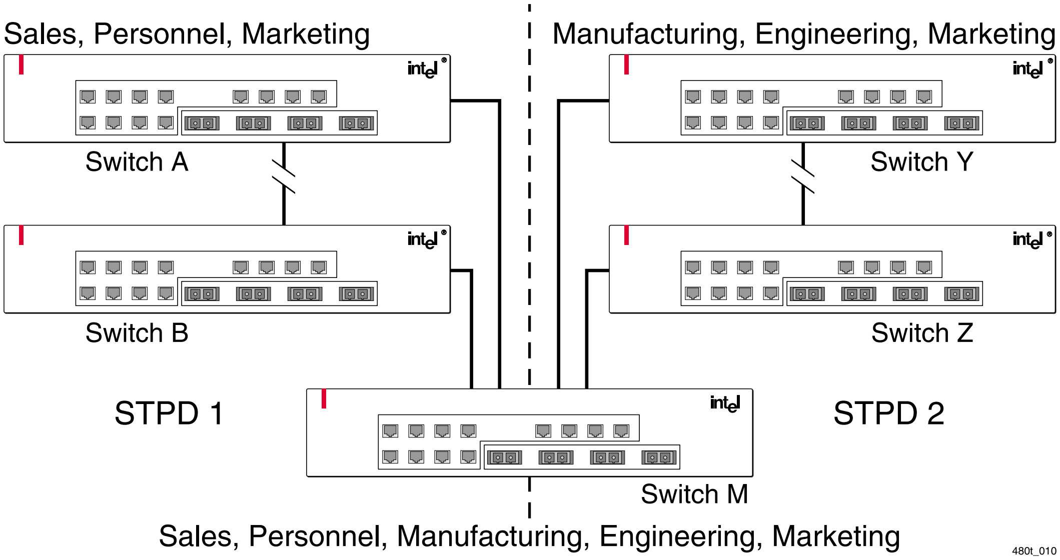
Sales, Personnel, Marketing
Manufacturing, Engineering, Marketing
Switch A
Switch Y
Switch B
Switch Z
STPD 1
STPD 2
Switch M
Sales, Personnel, Manufacturing, Engineering, Marketing
480t_010

Figure 9.1: Multiple Spanning Tree Domains - VLAN tagging for trunk connections

When the switches in this configuration start, STP configures each STPD such that there are no active loops in the topology. STP can configure the topology in several ways to make it loop-free.

In Figure 9.1, the connection between Switch A and Switch B is put into blocking state, and the connection between Switch Y and Switch Z is put into blocking state. After STP converges, all the VLANs can communicate, and all bridging loops are prevented.

The VLAN ***Marketing***, which was not assigned to either STPD1 or STPD2, communicates using all five switches. The topology has no loops, because STP has already blocked the port connection between Switches A and B, and between Switches Y and Z.

Be careful when configuring your VLANs within a single STPD. Figure 9.2 illustrates an incorrect network configuration using a single STPD.

The STP configuration disables the ability of the switches to forward VLAN traffic.

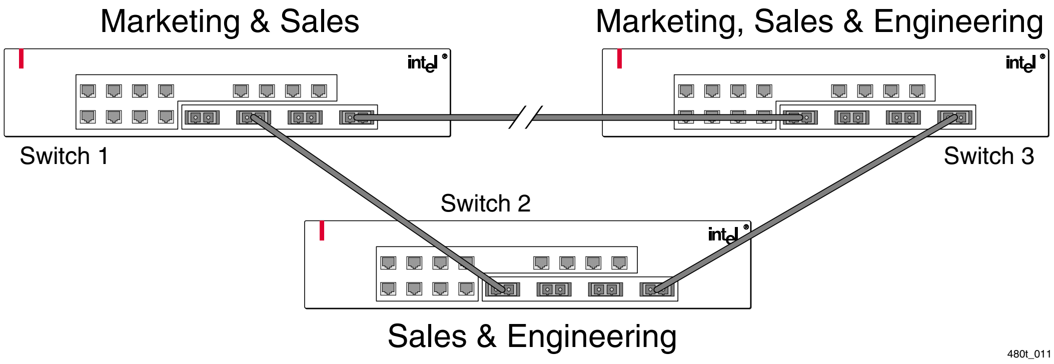

Figure 9.2: Tag-based STP configuration -Incorrect

The tag-based network in Figure 9.2 has this configuration:

- Switch 1 contains VLAN ***Marketing*** and VLAN ***Sales***.
- Switch 2 contains VLAN ***Engineering*** and VLAN ***Sales***.
- Switch 3 contains VLAN ***Marketing***, VLAN ***Engineering***, and VLAN ***Sales***.
- The tagged trunk connections for three switches form a triangular loop that is not permitted in an STP topology.
- All VLANs in each switch are members of the same STPD.

STP may block traffic between Switch 1 and Switch 3 by disabling the trunk ports for that connection on each switch.

Switch 2 has no ports assigned to VLAN marketing. Therefore, if the trunk for VLAN marketing on Switches 1 and 3 is blocked, the traffic for VLAN marketing will not be able to traverse the switches.

Configuring STP

We recommend that you do not configure STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

To configure STP:

STPD, VLAN, and QoS profile names must be unique. For example, a name used to identify a VLAN cannot be used for an STPD or a QoS profile.

1. Create one or more STP domains using this command:

   ```
   create stpd <stpd_name>
   ```

2. Add one or more VLANs to the STPD using this command:

   ```
   configure stpd <stpd_name> add vlan <name>
   ```

3. Enable STP for one or more STP domains using this command:

   ```
   enable stpd {<stpd_name>}
   ```

All VLANs belong to an STPD. If you do not want to run STP on a VLAN, you must add the VLAN to an STPD that is disabled.

After you create the STPD, you can optionally configure STP parameters for the STPD.

You can configure these parameters on each STPD:

- Hello time (default value 2 seconds)
- Forward delay (default value 15 seconds)
- Max age (default value 20 seconds)
- Bridge priority (default value 32768)

You can configure these parameters on each port:

- Path cost
- Port priority

The device supports the RFC 1493 Bridge MIB. You can only access the parameters of the s0 default STPD through this MIB.

Table 9.3 lists the commands used to configure STP. Press the Tab key in the command line interface for further command options.

Table 9.3: STP Configuration Commands

Command	Description
configure stpd <stpd_name> add vlan <name>	Adds a VLAN to the STPD.
configure stpd <stpd_name> forwarddelay <value>	Specifies the time (in seconds) that the ports in this STPD spend in the listening and learning states when the switch is the root bridge. The range is 4 through 30. Default setting is `15`.
configure stpd <stpd_name> hellotime <value>	Specifies the time delay (in seconds) between the transmission of BPDUs from this STPD when it is the root bridge. The range is 1 through 10. Default setting is 2.
configure stpd <stpd_name> maxage <value>	Specifies maximum age of BPDU in this STPD. The range is 6 through 40. Default setting is 20. Note that the time must be greater than, or equal to 2 x (Hello Time + 1) and less than, or equal to 2 x (Forward Delay –1).
configure stpd <stpd_name> port cost <value> <portlist>	Specifies the path cost of the port in this STPD. The range is 1 through 65,535. The switch automatically assigns a default path cost based on the speed of the port, as follows: • For a 100-Mbps port, the default cost is 19. • For a 1000-Mbps port, the default cost is 4.
configure stpd <stpd_name> port priority <value> <portlist>	Specifies the priority of the port in this STPD. By changing the port priority, you can make it more or less likely to become the Root Port. The range is 0 through 31. Default setting is 16. A setting of 0 indicates the lowest priority.
configure stpd <stpd_name> priority <value>	Specifies the priority of the STPD. By changing the priority, you can make it more or less likely to become the root bridge. • The range is 0 through 65,535. • The default setting is 32,768. • A setting of 0 indicates the highest priority.

Table 9.3: STP Configuration Commands (continued)

Command	Description
create stpd <stpd_name>	Creates an STPD. When created, an STPD has these default parameters: • Bridge priority—32,768 • Hello time—2 seconds • Forward delay—15 seconds
enable ignore-stp vlan <name>	Configures the switch to ignore the STP protocol, and not block traffic for the VLAN(s). This command is useful when multiple VLANs share the same physical ports, but only some of the VLANs require STP protection. The default setting is disabled.
enable stpd {<stpd_name>}	Enables the STP protocol for one or all STPDs. The default setting is disabled.
enable stpd <stpd_name> port {<portlist>}	Enables the STP protocol on one or more ports. If STPD is enabled for a port, Bridge Protocol Data Units (BPDUs) are generated on that port (if STP is enabled for the associated STPD). The default setting is enabled.
enable ignore-bpdu vlan <vlan-name>	Configures the switch to ignore the BPDU protocol on a VLAN.

STP Configuration Example

This example creates and enables an STPD named ***Backbone_st***. It assigns the ***Manufacturing*** VLAN to the STPD. It disables STP on ports 1 through 7, and port 12.

```
create stpd backbone_st
configure stpd backbone_st add vlan manufacturing
enable stpd backbone_st
disable stpd backbone_st port 1-7,12
```

Displaying STP Settings

To display STP settings, use this command:

```
show stpd {<stpd_name>}
```

This command displays:

- STPD name
- Bridge ID
- STPD configuration information

To display the STP state of a port, use this command:

```
show stpd <stpd_name> port <portlist>
```

This command displays:

- STPD port configuration
- STPD state (Root Bridge, and so on)
- STPD port state (forwarding, blocking, and so on)

Disabling and Resetting STP

To disable STP or return STP settings to their defaults, use the commands listed in Table 9.4. For further command options, press the Tab key in the command line interface.

Table 9.4: STP Disable and Reset Commands

Command	Description
delete stpd <stpd_name>	Removes an STPD. An STPD can only be removed if all VLANs were deleted from it. The default STPD, `s0`, cannot be deleted.
disable ignore-stp vlan <name>	Allows a VLAN to use STP port information.
disable stpd [<stpd_name> | all]	Disables the STP mechanism on a particular STPD or for all STPDs.
disable ignore-bpdu vlan <vlan-name>	Disables the ignoring of Bridge Protocol Data Units (BPDUs) on a VLAN.
disable stpd port <portlist>	Disables STP on one or more ports. Disabling STP on one or more ports puts those ports in forwarding state; all BPDUs received on those ports are disregarded.
unconfigure stpd {<stpd_name>}	Restores default STP values to a particular STPD or to all STPDs.

10 Quality of Service (QoS)

This chapter describes the concept of Quality of Service (QoS) and explains how to configure QoS on the Intel® NetStructure™ 480T routing switch.

Overview of Policy-Based Quality of Service

Policy-based QoS allows you to assign specific levels of service to different traffic types traversing the switch. Policy-based QoS is an effective control mechanism for networks that have heterogeneous traffic patterns.

Using Policy-based QoS, you can specify the service level for a particular traffic type.

Policy-based QoS allows you to protect bandwidth for important categories of applications or specifically limit the bandwidth associated with less critical traffic.

For example, if voice-over IP traffic requires a reserved amount of bandwidth to function properly, using policy-based QoS, you can reserve sufficient bandwidth to preserve latency characteristics critical to this type of application. Less critical applications can also be limited to preserve bandwidth.

The switch contains separate hardware queues on every physical port. Each hardware queue is programmed with bandwidth-management and prioritization parameters. The bandwidth-management and prioritization parameters that modify the forwarding behavior of the switch affect how the switch transmits traffic for a given hardware queue on a physical port.

The switch tracks and enforces the minimum and maximum percentage utilization transmitted on every hardware queue for every port.

Prioritization is utilized when two or more hardware queues on the same physical port are contending for transmission, as long as their respective bandwidth-management parameters have been satisfied.

Random Early Detection

Policy-based QoS can be configured to perform per-port Random Early Detection (RED) and drop-probability. Using this capability, the switch detects when traffic is filling up in any of the hardware queues, and performs a random discard on subsequent packets, based on the configured RED drop-probability.

Instead of dropping sessions during times when the queue depth is exceeded, RED causes the switch to lower session throughput. The destination node detects the dropped packet, and, using TCP windowing mechanisms, slows the transmission from the source node. RED drop-probability is configured on a system-wide basis, and has a valid range from 0% to 100%.

Policy-Based Routing and Route Load Sharing

For information on policy-based routing and route load sharing (link aggregation) refer to "Policy-Based Routing and Route Load-Sharing" on page 190.

Performance Impact

Utilizing any aspect of policy-based Quality of Service has zero impact on switch performance. Using even the most complex traffic groupings is costless in terms of switch performance.

Applications and Types of QoS

Applications vary significantly in QoS requirements. These applications are ones that you will most commonly encounter and need to prioritize:

- Voice applications
- Video applications
- Critical database applications
- Web browsing applications
- File server applications

General guidelines for each traffic type are given below and summarized in Table 10.1 on page 139. These are general guidelines and not strict recommendations.

After the QoS parameters are set, you can monitor the performance of the application to determine if the actual behavior of the applications matches your expectations. It is important to understand the needs and behavior of the particular applications you want to protect or limit. Behavioral aspects to consider include:

- Bandwidth needs
- Sensitivity to latency and jitter
- Sensitivity and impact of packet loss

Voice Applications

Voice applications typically demand small amounts of bandwidth. However, the bandwidth must be constant and predictable because these applications are typically sensitive to latency (inter-packet delay) and jitter (variation in inter-packet delay).

The most important QoS parameter to establish for voice applications is minimum bandwidth, followed by priority.

Video Applications

Video applications are similar in QoS needs to voice applications, with the exception that bandwidth requirements are somewhat larger, depending on the encoding. It is important to understand the behavior of the video application being used.

For example, in the playback of stored video streams, some applications can transmit large amounts of data for multiple streams in one spike, with the expectation that the endstations will buffer significant amounts of video-stream data. This can present a problem to the network infrastructure, because it must be capable of buffering the transmitted spikes where there are speed differences (for example, going from Gigabit Ethernet to Fast Ethernet).

The key QoS parameters for video applications include:

- Minimum bandwidth
- Priority
- Buffering (depending on the behavior of the application)

Critical Database Applications

Database applications, such as those associated with ERP (enterprise resource planning), typically do not demand significant bandwidth and are tolerant of delay. You can establish a minimum bandwidth using a priority less than that of delay-sensitive applications.

Web Browsing Applications

Use full-duplex links when deploying policy-based QoS. Half-duplex operation on links can make delivery of guaranteed minimum bandwidth impossible.

QoS needs for Web-browsing applications cannot be generalized into a single category. For example, ERP applications that use a browser front-end may be more important than those retrieving daily news information. Traffic groupings can be distinguished from each other by their server source and destinations.

Most browser-based applications are distinguished by asymmetric data flow (small data flows from the browser client, large data flows from the server to the browser client). An exception to this can be created by some Java§ applications. In addition, Web-based applications are generally tolerant of latency, jitter, and some packet loss. However small, packet-loss can have a large impact on perceived performance due to the nature of TCP.

The relevant parameter for protecting browser applications is minimum bandwidth. The relevant parameter for preventing non-critical browser applications from overwhelming the network is maximum bandwidth. In addition, RED can be used to reduce session loss if the queue that floods Web traffic becomes over-subscribed.

File Server Applications

File serving typically poses the greatest demand on bandwidth, although file server applications are tolerant of latency, jitter, and some packet loss, depending on the network operating system and the use of TCP or UDP.

Table 10.1: Traffic Type and QoS Guidelines

Traffic Type	Key QoS Parameters
Voice	Minimum bandwidth, priority
Video	Minimum bandwidth, priority, buffering (varies)
Database	Minimum bandwidth
Web browsing	Minimum bandwidth for critical applications, maximum bandwidth for non-critical applications, RED
File server	Minimum bandwidth

Building Blocks

The service that a particular type of traffic or traffic grouping receives is determined by assigning that traffic to a QoS profile. A QoS profile is characterized by minimum and maximum bandwidth and prioritization settings that define a desired class of service.

Assigning QoS Attributes

To assign QoS attributes you must define three interrelated QoS building blocks in three steps:

1. Define a QoS profile.

   QoS profile—A class of service that is defined through minimum and maximum bandwidth parameters, configuration of buffering and RED, and prioritization settings. The bandwidth and level of service that a particular type of traffic or traffic grouping receives is determined by assigning it to a QoS profile.

2. Assign one or more traffic groupings to a QoS profile to create a QoS policy.

   Traffic grouping—A classification or traffic type that has one or more attributes in common. These can range from a physical port to a VLAN to IP Layer 4 port information. Traffic groupings are assigned to QoS profiles to modify switch forwarding behavior. Traffic groupings transmitting out the same port that are assigned to a particular QoS profile share the assigned bandwidth and prioritization characteristics, and hence share the class of service.

   QoS policy—The combination that results from assigning a traffic grouping to a QoS profile.

3. Monitor the performance of the application with the QoS monitor to determine whether the policies are meeting the desired results.

QoS Profiles

Eight default QoS profiles are provided. The default QoS profiles cannot be deleted. QoS profiles are linked to hardware queues. There are multiple hardware queues per physical port. The default QoS profile names and their queues are described in Table 10.2.

Table 10.2: Default QoS Profile Names and Queues

QoS Profile Name	Hardware Queue
Qp1	Q0
Qp2	Q1
Qp3	Q2
Qp4	Q3
Qp5	Q4
Qp6	Q5
Qp7	Q6
Qp8	Q7

Each physical port contains all of the hardware queues listed in Table 10.2.

The parameters that make up a QoS profile include:

- **Minimum bandwidth** – The minimum percentage of total link bandwidth that should be reserved for use by a hardware queue on a physical port. Bandwidth unused by that queue may be used by other queues. The minimum bandwidth for all queues should add up to less than 90%. The default value on all minimum bandwidth parameters is 0%.

- **Maximum bandwidth** – The maximum percentage of total link bandwidth that may be transmitted by a hardware queue on a physical port. The default value on all maximum bandwidth parameters is 100%.
- **Priority** – The level of priority assigned to a hardware queue on a physical port. The switch has eight different available priority settings. By default, each of the default QoS profiles is assigned a unique priority. Prioritization is used under these circumstances:
  - When two or more hardware queues on the same physical port are contending for transmission on the same physical port, only after their respective bandwidth-management parameters have been satisfied. If two hardware queues on the same physical port have the same priority, a round-robin algorithm is used for transmission.
  - When configured to do so, the priority of a QoS profile determines the 802.1p bits used in the priority field of a transmitted packet.

See "Configuring DiffServ" on page 151

  - The priority of a QoS profile determines the DiffServ code point value used in an IP packet when the packet is transmitted.
- **Buffer** – This parameter reserves buffer memory for use exclusively by a QoS profile across all affected ports. The default value for buffer settings is 0%. The *sumvalue* of all QoS profile buffer parameters should not exceed 100%. Reserving buffer memory for a QoS profile affects the dynamic buffer space available to other QoS profiles. You should not modify the buffer parameter unless specific situations and application behavior indicates this need.

A QoS profile does not alter the behavior of the switch until it is assigned to a traffic grouping. Remember that QoS profiles are linked to hardware queues. There are multiple hardware queues per physical port. By default, a QoS profile links to the identical hardware queue across all the physical ports of the switch.

The settings for the default QoS profiles are summarized in Table 10.3. For further command options, press the Tab key in the command line interface.

Table 10.3: Default QoS Profiles

Profile Name	Hardware Queue	Priority	Buffer	Minimum Bandwidth	Maximum Bandwidth
Qp1	Q0	Low	0	0%	100%
Qp2	Q1	Lowhi	0	0%	100%
Qp3	Q2	Normal	0	0%	100%
Qp4	Q3	Normalhi	0	0%	100%
Qp5	Q4	Medium	0	0%	100%
Qp6	Q5	Mediumhi	0	0%	100%
Qp7	Q6	High	0	0%	100%
Qp8	Q7	Highhi	0	0%	100%

Configuring a QoS Profile

Table 10.4 lists the commands used to configure QoS. For further command options, press the Tab key in the command line interface.

Table 10.4: QoS Configuration Commands

Command	Description
configure qosprofile <qosprofile> {minbw <percent>} {maxbw <percent>} {priority <level>} {<portlist> | maxbuf <percent> minbuf <percent> [K | M]}	Configures a QoS profile. Specify: • `minbw`—The minimum buffer percentage guaranteed to be available to this queue for transmission. The default setting is 0. • `maxbw`—The maximum buffer percentage this queue is permitted to use for transmission. The default setting is `100`. • `priority`—The service priority for this queue. Settings include low, normal, medium, and high. The default setting is low. Available only in egress mode. • `maxbuf`—The maximum buffer for each queue, keeps a single queue from using all un-allocated buffer space. • `minbuf`—The minimum buffer for each queue. • `K/M`—Specifies kilobytes or megabytes with respect to buffer size.
configure ports [all | mgmt | <portnumber>] qosprofile <qosprofile>	Configures one or more ports to use a particular QoS profile. Available only in ingress mode.
configure red drop-probability <percent>	Configures the Random Early Detect (RED) drop-probability. The percentage range is 0 to 100.
configure vlan <name> qosprofile [<qosprofile> | none]	Configures a VLAN to use a particular QoS profile.
disable red ports	Disables RED on one or all ports.
enable red port <portnumber>	Enables RED on a port.

Modifying a QoS Profile

You can modify the default profiles as desired. To modify the parameters of an existing QoS profile, use this command:

```
configure qosprofile <qosprofile> {minbw <percent>}
{maxbw <percent>} {priority <level>} {minbuf
<percent>} {maxBuf <percent>} [K | M]
```

Traffic Groupings and Creating a QoS Policy

Use full-duplex links when deploying policy-based QoS. Half-duplex operation on links can make delivery of guaranteed minimum bandwidth impossible.

After a QoS profile is modified for bandwidth and priority, you assign the profile to a particular traffic grouping. A QoS profile is assigned to a specific traffic grouping to create a QoS policy. A traffic grouping is a classification of traffic that has one or more attributes in common.

Traffic groupings can be separated into these categories:

- IP-based information, such as IP source/destination and TCP/UDP port information
- Destination MAC (MAC QoS groupings)
- Explicit packet class of service information, such as 802.1p or DiffServ (IP TOS)
- Physical/logical configuration (physical source port or VLAN association)

If a given packet matches two or more grouping criteria, there is a predetermined precedence for which traffic grouping will apply. In general, the more specific traffic grouping takes precedence.

By default, all traffic groupings are placed in the QoS profile Qp1.The supported traffic groupings and their options by QoS mode are listed in Table 10.5. The groupings are listed in order of precedence (highest to lowest).

Table 10.5: Traffic Groupings by QoS Mode

IP Information (Access Lists) Groupings
• Access list precedence determined by user configuration
• IP destination

Table 10.5: Traffic Groupings by QoS Mode (continued)

IP Information (Access Lists) Groupings
Destination Address MAC-based Groupings • Permanent • Dynamic • Blackhole • Broadcast/unknown rate limiting
Explicit Packet Class of Service Groupings • DiffServ (IP TOS) • 802.1p
Physical/Logical Groupings • Source port • VLAN

IP-Based Traffic Groupings

IP-based traffic groupings are based on any combination of:

- IP source or destination address
- TCP/UDP or other Layer 4 protocol
- TCP/UDP port information

IP-based traffic groupings are defined using access lists (see chapter 16). By supplying a named QoS profile at the end of the access list command syntax, you can prescribe the bandwidth-management and priority handling for that traffic grouping. This level of packet filtering has no negative impact on performance.

MAC-Based Traffic Groupings

You can assign QoS profiles to destination MAC addresses. MAC-based traffic groupings are configured using this command:

```
create fdbentry <MAC address> vlan <vlan>
[blackhole | port <portlist> | dynamic qosprofile
<qosprofile>]
```

The MAC address options are:

- Permanent
- Dynamic
- Blackhole
- Broadcast/unknown rate limiting

Permanent MAC Addresses

Permanent MAC addresses can be assigned a QoS profile whenever traffic is destined for the MAC address. You can do this when you create a permanent FDB entry. For example:

```
create fdbentry 00:11:22:33:44:55 vlan default port
4 qosprofile qp2
```

Dynamic MAC Addresses

Dynamic MAC addresses can be assigned a QoS profile whenever traffic is destined for the MAC address. For any port on which the specified MAC address is learned in the specified VLAN, the port is assigned the specified QoS profile. For example:

```
create fdbentry 00:11:22:33:44:55 vlan default
dynamic qosprofile qp3
```

The QoS profile is assigned when the MAC address is learned. When a client location changes, the assigned QoS profile moves with the device. If the MAC address entry already exists in the FDB, you can clear the forwarding database so that the QoS profile can be applied when the entry is added again. The command to clear the FDB is:

```
clear fdb
```

Blackhole MAC Address

Using the **blackhole** option configures the switch to not forward any packets to the destination MAC address on any ports for the VLAN specified. The **blackhole** option is configured using this command:

```
create fdbentry 00:11:22:33:44:55 vlan default
blackhole
```

Broadcast/Unknown Rate Limiting MAC Address

IP multicast traffic is subject to broadcast and unknown rate limiting only when IGMP snooping is disabled. Refer to "IGMP Snooping" on page 278.

It is possible to assign broadcast and unknown destination packets to a QoS profile that has the desired priority and bandwidth parameters. Broadcast/unknown rate limiting is an extension of the QoS feature used for destination MAC addresses.

For example, if you want to limit broadcast and unknown traffic on the VLAN ***default*** to the bandwidth and priority defined in QoS profile ***qp3***, the command is:

```
create fdbentry ff:ff:ff:ff:ff:ff vlan default
dynamic qp3
```

Verifying MAC-Based QoS Settings

To verify any of the MAC-based QoS settings, use either of these two commands:

```
show fdb perm
```

```
show qosprofile <qosprofile>
```

Explicit Class of Service Traffic Groupings (802.1p and DiffServ)

This category of traffic groupings describes what is sometimes referred to as explicit packet marking, and refers to information contained within a packet that is intended to explicitly determine a class of service. This includes:

- IP Differentiated Services (DiffServ) code points, also known as IP TOS bits
- Prioritization bits used in IEEE 802.1p packets

An advantage of explicit packet marking is that the class of service information can be carried throughout the network infrastructure, without repeating what may be complex traffic grouping policies at each switch location. Another advantage is that endstations can perform their own packet marking on an application-specific basis. The 480T routing switch can observe and manipulate packet marking information with no performance penalty.

The documented capabilities for 802.1p priority markings or DiffServ capabilities are not impacted by the switching or routing configuration of the switch. For example, 802.1p information can be preserved across a routed switch boundary and DiffServ code points can be observed or overwritten across a Layer 2 switch boundary.

Configuring 802.1p Priority

The switch supports the standard 802.1p priority bits that are part of a tagged Ethernet packet. The 802.1p bits can be used to prioritize the packet, and assign it to a particular QoS profile.

When a packet arrives at the switch, the 802.1p priority field is examined, and can be mapped to a specific hardware queue for subsequent transmission. The 802.1p priority field is located directly following the 802.1Q type field, and preceding the 802.1Q VLAN ID, as shown in Figure 10.1.

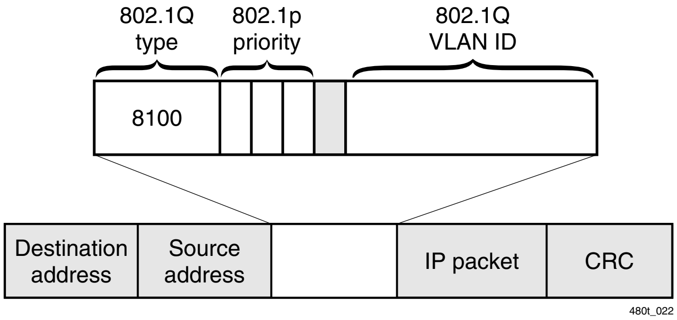

Figure 10.1: Ethernet packet encapsulation

Observing 802.1p Information

When ingress traffic that contains 802.1p prioritization information is detected by the switch, the traffic is mapped to various hardware queues on the egress port of the switch. The 480T routing switch supports eight hardware queues. The hardware queues determine the bandwidth-management and priority characteristics used when transmitting packets.

To control the mapping of 802.1p prioritization values to hardware queues, 802.1p prioritization values can be mapped to a QoS profile. The default mapping of each 802.1p priority value to QoS profile is described in Table 10.6.

Table 10.6: 802.1p Priority Value-to-QoS Profile Mapping

Priority Value	QoS Profile
0	Qp1
1	Qp2
2	Qp3
3	Qp4
4	Qp5
5	Qp6
6	Qp7
7	Qp8

As described in Table 10.2, by default a QoS profile is mapped to a hardware queue, and each QoS profile has configurable bandwidth parameters and priority. In this way, an 802.1p priority value detected on ingress can be mapped to a particular QoS profile with specified bandwidth-management and priority behavior.

To change the default mappings of QoS profiles to 802.1p priority values, use the command:

```
configure dot1p ethertype <dot1p_priority>
```

Replacing 802.1p Priority Information

By default, 802.1p priority information is not replaced or manipulated, and the information observed on ingress is preserved when transmitting the packet. This behavior is not affected by the switching or routing configuration of the switch.

However, the switch is capable of inserting and/or overwriting 802.1p priority information when it transmits an 802.1Q tagged frame. If 802.1p replacement is enabled, the 802.1p priority information that is transmitted is determined by the hardware queue that is used when transmitting the packet. To replace 802.1p priority information, use the command:

```
enable dot1p replacement ports [<portlist> | all]
```

802.1p priority information is replaced according to the hardware queue that is used when transmitting from the switch. The mapping is described in Table 10.7. This mapping cannot be changed.

Table 10.7: 802.1p Priority Value-to-Hardware Queue Mapping

Hardware Queue	802.1p Priority Value
Q0	0
Q1	1
Q2	2
Q3	3
Q4	4
Q5	5
Q6	6
Q7	7

802.1p Commands

Table 10.8 shows the commands used to configure 802.1p priority. For further command options, press the Tab key in the command line interface.

Table 10.8: 802.1p Configuration Commands

Command	Description
configure dot1p ethertype <dot1p_priority>	Configures the default QoS profile to 802.1p priority mapping. The value for dot1p_priority is an integer between 0 and 7.

Table 10.8: 802.1p Configuration Commands (continued)

Command	Description
disable dot1p replacement ports [<portlist> | all]	Disables the ability to overwrite 802.1p priority values for a given set of ports.
enable dot1p replacement ports [<portlist> | all]	Enables the 802.1p priority field to be overwritten on egress according to the QoS profile to 802.1p priority mapping for a given set of ports.
show dot1p	Displays the 802.1p-to-QoS profile mappings.

Configuring DiffServ

Contained in the header of every IP packet is a field for IP Type of Service (TOS), also referred to as the DiffServ field. The DiffServ or TOS field is used by the switch to determine the type of service provided to the packet. Figure 10.2 shows the encapsulation of an IP packet header.

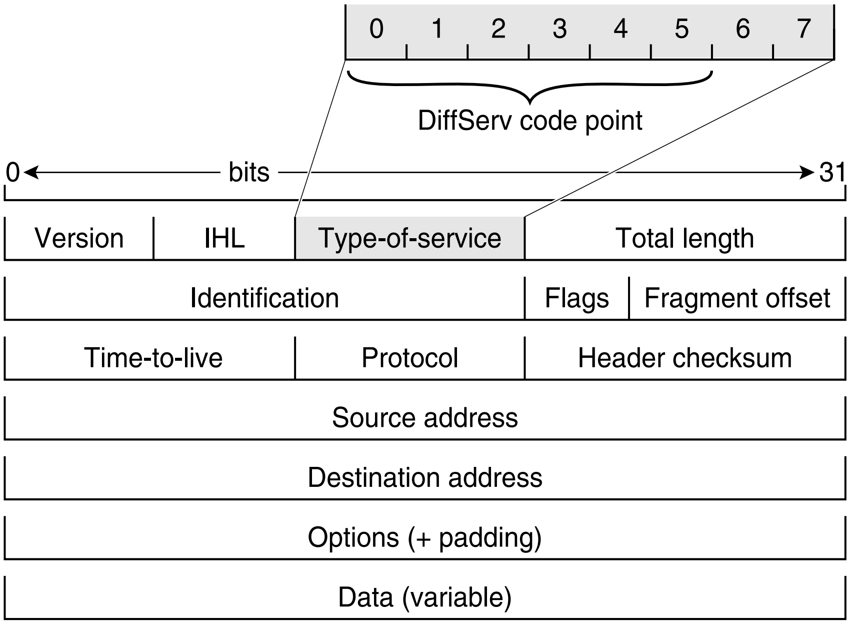

Figure 10.2: IP packet header encapsulation

Observing DiffServ Information

When a packet arrives at the switch on an ingress port, the switch examines the first six of eight TOS bits. These bits are called the code point.

The switch can assign the QoS profile used to subsequently transmit the packet based on the code point. The QoS profile controls a hardware queue used when transmitting the packet out of the switch, and determines the forwarding characteristics of a particular code point.

You can enable or disable the observance of DiffServ information. By default it is disabled. To enable observance of DiffServ information use the command:

```
enable diffserv examination ports [<portlist> |
all]
```

Changing DiffServ Code Point Assignments in the QoS Profile

Because the code point uses six bits, it has 64 possible values (2^6 = 64). By default, the values are grouped and assigned to the default QoS profiles listed in Table 10.9.

Table 10.9: Default Code Point-to-QoS Profile Mapping

Code Point	QoS Profile
0-7	Qp1
8-15	Qp2
16-23	Qp3
24-31	Qp4
32-39	Qp5
40-47	Qp6
48-55	Qp7
56-63	Qp8

You can change the QoS profile assignment for all 64 code points. Use this command:

```
configure diffserv examination code-point <code-
point> qosprofile <qosprofile> ports [<portlist>]
```

Once assigned, the rest of the switches in the network prioritize the packet using the characteristics specified by the QoS profile.

Replacing DiffServ Code Points

You can configure the switch to change the DiffServ code point in the packet prior to the packet being transmitted by the switch. This is accomplished with no impact on switch performance.

The DiffServ code point value used in overwriting a packet is determined by the 802.1p priority value. As described in the section “Replacing 802.1p Priority Information,” the 802.1p priority value is, in turn, determined by the hardware queue used when transmitting a packet.

It is not necessary to receive or transmit 802.1Q tagged frames, only to understand that the egress hardware queue, which also determines the 802.1p priority value, can also be configured to determine the DiffServ value if you want to replace the DiffServ code points.

To enable the replacement of DiffServ code points you must enable both 802.1p replacement and DiffServ replacement using these commands:

```
enable dot1p replacement ports [<portlist> | all]
```

```
enable diffserv replacement ports [<portlist> |
all]
```

The default 802.1p priority value to code point mapping is described in Table 10.10.

Table 10.10: Default 802.1p Priority Value-to-Code Point Mapping

Hardware Queue	802.1p Priority Value	Code Point
0	0	0
1	1	8
2	2	16
3	3	24
4	4	32
5	5	40
6	6	48
7	7	56

You can change the 802.1p priority to DiffServ code point mapping to any code point value using this command:

```
configure diffserv replacement priority vpri
<number> code-point <code-point> ports [<portlist>]
```

By doing so, the hardware queue used to transmit a packet determines the DiffServ value replaced in the IP packet.

To verify the DiffServ configuration, use the command:

```
show ports <portlist> info {detail}
```

Table 10.11 describes the commands used to configure DiffServ. For further command options, press the Tab key in the command line interface.

Table 10.11: DiffServ Configuration Commands

Command	Description
configure diffserv examination code-point <code-point> qosprofile <qosprofile> ports [<portlist>]	Configures the default ingress DiffServ code points to QoS profile mapping. The `<code-point>` is a 6-bit value in the IP-TOS byte in the IP header. You can specify up to 64 different code points for each port.
configure diffserv replacement priority vpri <number> code-point <code-point> ports [<portlist>]	Configures the default egress DiffServ replacement mapping.
disable diffserv examination ports [<portlist> | all]	Disables the examination of the DiffServ field in an IP packet.
disable diffserv replacement ports [<portlist> | all]	Disables the replacement of DiffServ code points in packets transmitted by the switch.
enable diffserv examination ports [<portlist> | all]	Enables the DiffServ field of an ingress IP packet to be examined by the switch in order to select a QoS profile. The default setting is disabled.
enable diffserv replacement ports [<portlist> | all]	Enables the DiffServ code point to be overwritten in packets transmitted by the switch. Eight user-defined code points can be configured on each port. The 802.1p priority bits (3-bits) are used to select one of the eight code points. The default setting is disabled.
unconfigure diffserv examination ports [<portlist>]	Removes the DiffServ examination code point from a port.
unconfigure diffserv replacement ports [<portlist>]	Removes the DiffServ replacement mapping from a port.

DiffServ Example

In this example, we use DiffServ to signal a class of service throughput and assign any traffic coming from network 10.1.2.x with a specific DiffServ code point. This allows all other network switches to send and observe the Diffserv code point instead of repeating the same QoS policy on every network switch.

Configure the switch that handles incoming traffic from network 10.1.2.x as follows:

1. Configure parameters of the QoS profile Qp3:

   ```
   configure qp3 min 10 max 100
   ```

2. Assign a traffic grouping for traffic from network 10.1.2.x to Qp3:

   ```
   create access-list TenOneTwo
   configure TenOneTwo 10.1.2.0/24 permit qp3
   ```

3. To enable the switch to overwrite the DiffServ code point:

   ```
   enable dot1p replacement ports all
   ```

   ```
   enable diffserv replacement ports all
   ```

4. Configure the switch so that other switches can signal the class of service that this switch should observe:

   ```
   enable diffserv examination ports all
   ```

Table 10.3 indicates that Qp3 is tied to hardware queue Q2. When replacement is enabled all traffic sent out Q2 will contain code point value 16 (according to Table 10.10). If this is the desired code point to use, all traffic from 10.1.2.x is sent out Qp3 (at 10% minimum and 100% maximum) with a code point value of 16.

Physical and Logical Groupings

Two traffic groupings exist in this category:

- Source port
- VLAN

Source Port

A source port traffic grouping implies that any traffic sourced from this physical port uses the indicated QoS profile when the traffic is transmitted out to any other port. To configure a source port traffic grouping, use this command:

```
configure ports [all | mgmt | <portnumber>]
qosprofile <qosprofile>
```

In the following example, all traffic sourced from port 7 uses the QoS profile named ***qp3*** when being transmitted.

```
configure ports 7 qosprofile qp3
```

VLAN

A VLAN traffic grouping indicates that all intra-VLAN switched traffic and all routed traffic sourced from the named VLAN uses the indicated QoS profile. To configure a VLAN traffic grouping, use this command:

```
configure vlan <name> qosprofile [<qosprofile> |
none]
```

For example, all devices on VLAN ***servnet*** require use of the QoS profile ***qp4***. The command to configure this example is:

```
configure vlan servnet qosprofile qp4
```

Verifying Physical and Logical Groupings

To verify settings on ports or VLANs, use the command:

```
show qosprofile <qosprofile>
```

The same information is also available for ports or VLANs using:

```
show ports info
```

or

```
show vlan
```

Verifying Configuration and Performance

You can use the information in this section to verify the QoS configuration and monitor the use of the QoS policies that are in place.

QoS Monitor

The QoS monitor is a utility that monitors the hardware queues associated with any port(s). The QoS monitor keeps track of the number of frames and the frames per second that a specific queue is responsible for transmitting on a physical port. Two options are available: a real-time display, and a separate option for retrieving information in the background and writing it to the log.

Table 10.12 describes the QoS monitor commands. For further command options, press the Tab key in the command line interface.

Table 10.12: QoS Monitor Commands

Command	Description
disable qosmonitor	Disables the QoS monitoring capability.
enable qosmonitor port [<port> | mgmt]	Enables the QoS monitoring capability on the switch. When no port is specified, the QoS monitor automatically samples all the ports. Error messages are logged to the syslog if the traffic exceeds the parameters of the QoS profile(s). The default setting is disabled.
show ports {<portlist>} qosmonitor	Displays real-time QoS statistics for one or more ports.

Real-Time Performance Monitoring

The real-time display scrolls through the given portlist to provide statistics. Screens for packet count and packets per second can be chosen. The particular port being monitored at that time is indicated by an asterisk (*) appearing after the port number in the display.

The command for real-time viewing is:

```
show ports {<portlist>} qosmonitor
```

QoS monitor sampling is configured as follows:

- The port is monitored for 20 seconds before the switch moves to the next port in the list.

- A port is sampled for five seconds before the packets per second (pps) value is displayed on the screen.

Background Performance Monitoring

Monitoring QoS in the background places the transmit counter and any overflow information into the switch log. The log notification appears if one of the queues experiences an overflow condition since the last time it was sampled.

An overflow entry indicates that a queue was over-subscribed at least temporarily, and is useful for determining correct QoS settings and potential over-subscription issues.

Displaying QoS Information

The QoS monitor can also be used to verify the QoS configuration and monitor the use of the QoS policies that are in place. To display QoS information on the switch, use this command:

```
show qosprofile <qosprofile>
```

Displayed information includes:

- QoS profile name
- Minimum bandwidth
- Maximum bandwidth
- Priority
- A list of all traffic groups to which the QoS profile is applied

Additionally, you can display QoS information from the traffic grouping perspective by using one or more of these applicable commands:

- To display destination MAC entries and their QoS profiles.

  ```
  show fdb permanent
  ```

- To display general switch information :

  ```
  show switch
  ```

- To display the QoS profile assignments to the VLAN.

  ```
  show vlan
  ```

- To display information including QoS information for the port.

```
show ports info {detail}
```

Modifying a QoS Policy

If you change the parameters of a QoS profile after a QoS policy was created (by applying a QoS profile to a traffic grouping), the timing of the configuration change depends on the traffic grouping involved. To have a change in QoS profile affect a change in the QoS policy, these rules apply:

- For destination MAC-based grouping (other than permanent), clear the MAC FDB using the command `clear fdb`. This command should also be issued after a policy is first formed, as the policy must be in place before an entry is made in the MAC FDB. For permanent destination MAC-based grouping, re-apply the QoS profile to the static FDB entry, as documented. You can also save and reboot the switch.
- For physical and logical groupings of a source port or VLAN, re-apply the QoS profile to the source port or VLAN, as documented. You can also save and reboot the switch.

QoS Profile Buffer

Although the QoS profile buffer can be set to a value greater than 100, the maximum effective setting is 100%.

Maximum QoS Buffer

The `maxbuf` parameter allows you to set a maximum buffer for each queue, so that a single queue will not consume all of the un-allocated buffer space.

The `maxbuf` values can be set in kilobit or megabit increments. The minimum value is zero K and the maximum is 16,384K. The default value is zero K. Unless you have explicit reasons, do not modify these parameters. Only unique situations require any non-default configurations of QoS.

To set the **`maxbuf`** value on a queue, use this command:

```
configure qosprofile <qos profile> minbw <percent>
maxbw <percent> priority <priority> maxbuf <number>
```

To view the **`maxbuf`** configuration, use this command:

```
show qosprofile
```

Bandwidth Settings and Their Impact

Bandwidth settings applied to QoS profiles used for ingress or egress traffic are expressed as a percentage of bandwidth. QoS profile bandwidth settings are in turn applied to queues on physical ports. The actual impact of the bandwidth setting is determined by the port speed (100 or 1000 Mbps) and by the actual granularity capabilities of the switch.

Maximum bandwidth settings

The maximum bandwidth percentage settings determine the port bandwidth available to each queue. Use Table 10.13 to determine the actual maximum bandwidth associated with each setting. If the maximum percentage bandwidth configured does not match one of the settings listed below, it is rounded up to the next setting.

Table 10.13: QoS Maximum Bandwidth Settings

Maximum Bandwidth Setting (%)	Maximum Bandwidth @ 100Mbps	Maximum Bandwidth @ 1000 Mbps
2%	2 Mbps	20 Mbps
3%	3.1 Mbps	30 Mbps
5%	4.9 Mbps	50 Mbps
7%	6.9 Mbps	69 Mbps
8%	7.9 Mbps	79 Mbps

Table 10.13: QoS Maximum Bandwidth Settings

Maximum Bandwidth Setting (%)	Maximum Bandwidth @ 100Mbps	Maximum Bandwidth @ 1000 Mbps
10%	9.6 Mbps	96 Mbps
11%	11.2 Mbps	112 Mbps
15%	15 Mbps	150 Mbps
20%	19 Mbps	190 Mbps
25%	25 Mbps	250 Mbps
30%	33Mbps	330 Mbps
35%	35 Mbps	350 Mbps
40%	42 Mbps	420 Mbps

Minimum bandwidth settings

The minimum bandwidth settings determine the reserved port bandwidth available to each queue. Table 10.14 shows actual reserved bandwidth for each setting. If the reserved percentage configured does not match the settings below, it is rounded up. If the actual bandwidth used is below the minimum bandwidth within a queue, other queues on that physical port can use it.

Table 10.14: QoS Profile Minimum Bandwidth

Minimum Bandwidth Setting (%)	Minimum Bandwidth@ 100 Mbps	Minimum Bandwidth @ 1000 Mbps
4%	4.2 Mbps	42 Mbps
6%	5.7 Mbps	57 Mbps

Table 10.14: QoS Profile Minimum Bandwidth

Minimum Bandwidth Setting (%)	Minimum Bandwidth@ 100 Mbps	Minimum Bandwidth @ 1000 Mbps
8%	7.5 Mbps	75 Mbps
9%	9.3 Mbps	93 Mbps
10%	10 Mbps	100 Mbps
20%	18.7 Mbps	187 Mbps
25%	26.3 Mbps	263 Mbps
35%	34 Mbps	340 Mbps
50%	49 Mbps	490 Mbps
60%	63 Mbps	630 Mbps
80%	79 Mbps	790 Mbps
89%	94 Mbps	940 Mbps

The sum of the minimum bandwidth values for the applied QoS profiles should be kept to less than 90% of available bandwidth.

If the minimum bandwidth settings exceed 90% it is possible, under a sustained situation of over-subscription, that a lower priority queue could become “starved” and not transmit traffic.

Bi-directional Rate Shaping for Layer 3 Routed VLANs

Bi-directional rate shaping allows you to perform bandwidth-management for Layer 2 and Layer 3 traffic flowing both to and from the switch.

You can utilize up to eight ingress rate-shaping queues per VLAN and eight egress rate-shaping queues per physical port. By defining a QoS profile's minimum and maximum bandwidth corresponding to the physical queue and port, you define committed information rates for each queue and port. Different bandwidth rates can be applied to ingress vs. egress traffic.

You can then create traffic groupings (e.g. physical port, VLAN, 1p, DiffServ, IP address, Layer 4 flow) for the eight pre-defined QoS profiles, thereby directing specific types of traffic to the desired queue. The traffic groupings used are not dependent on whether the traffic is switched or routed.

When you configure switch ports in L2 mode, MAC-block conflicts will not return error messages if L3 mode is later enabled.

The switch returns error messages on MAC-block conflicts when you add rate-shaped ports to VLANs.

MAC-block restrictions do not exist when using the switch as Layer 2 only.

Configuring Bi-Directional Rate Shaping

For bi-directional rate shaping to work, each VLAN requires a loopback port. This operates by directing all traffic from rate-shaped ports through the loopback port for that VLAN. To rate-shape ingress traffic, configure QoS normally on the loopback port for the VLAN.

The maximum bandwidth and traffic grouping defined in the QoS profile for the loopback port sets the rate limit for ingress traffic on rate-shaped ports in that VLAN.

Use these guidelines for bi-directional ingress rate shaping:

- You must configure a loopback port before adding rate-shaped ports to the VLAN.
- A loopback port cannot be used by an external device.

A loopback port cannot be used by an external device.

- A loopback port must have a unique loopback VLAN tag ID.
- Ingress traffic on a port that is configured to use the loopback port is rate-shaped.
- Ingress traffic on a port that is not configured to use the loopback port will not be rate-shaped.

- Unicast traffic from a non-rate-shaped port to a rate-shaped port within the VLAN will not be rate-shaped.
- The aggregate forwarding bandwidth of all rate-shaped ports in a VLAN is determined by the traffic groupings and bandwidth settings for the QoS profiles of the loopback port.

For egress rate shaping, simply set the maximum bandwidth of the QoS profile on the egress port.

Bi-Directional Rate Shaping Limitations

Consider these limitations when configuring bi-directional rate shaping:

- When configuring VLAN memberships, delete all rate-shaped ports before deleting the loopback port.
- If rate-shaped ports within a VLAN use different bandwidth parameters, set the priority of the QoS profiles on the loopback port and rate-shaped ports to low.
- Layer 2 switched rate-shaping only affects a single VLAN.
- IP forwarding must be enabled on the VLAN prior to adding the loopback port to a VLAN for L2 rate shaping.
- Ports that are tagged cannot be used for rate shaping.

If you have IP routing enabled and you do add a rate-shaped port to a VLAN, and the rate-shaped port is in the same port block as loopback or normal ports, the switch will return one of these error messages:

```
ERROR: Rate shaped port can’t be in the same block
as loopback port
```

```
ERROR: Normal port (port #) cannot share the block
with rate shaped port
```

Bi-Directional Rate Shaping Commands

To add the loopback port to the VLAN, use the command:

```
configure vlan <vlan name> add port <port> loopback-
vid <vlan_tag>
```

To enable the loopback port, use the following command:

```
restart port <loopback_port>
```

To add rate-shaped ports to the VLAN, use the following command:

```
configure vlan <vlan name> add port <portlist>
{tagged | untagged} {nobroadcast} soft-rate-limit
```

To delete rate-shaped ports from the VLAN, use the command:

```
configure vlan <vlan name> delete port <portlist>
```

To configure the rate-shaping parameters of the loopback port, use the normal QoS profile configuration command, as follows:

```
configure qosprofile <qosprofile> {minbw <pcnt>}
{maxbw <pcnt>} priority <level> {buffer <pcnt>}
{<portlist>} <loopback port number>
```

To display the bi-directional rate shaping configuration, use the command:

```
show vlan {<vlan name> | detail}
```

This command designates rate-shaped ports with an R and loopback ports with an L next to the port number.

To set the port speed of a loopback port, use the normal port configuration command, as follows:

```
configure ports <portlist> auto off {speed [ 100 |
1000]} duplex [half | full]
```

11 Enterprise Standby Router Protocol (ESRP)

Overview

We recommend that all switches using ESRP use the same version of firmware for interoperability. See "Software Upgrade and Boot Options" on page 419.

Enterprise Standby Router Protocol (ESRP) allows multiple switches to provide redundant routing services to users. From the workstation's perspective, there is only one default router that has one IP address and one MAC address, so ARP cache entries in client workstations do not need to be refreshed or aged-out.

Along with providing Layer 3 routing redundancy for IP and IPX§, ESRP also provides for Layer 2 redundancy. You can use these layered redundancy features in combination or independently. You do not have to configure the Intel® NetStructure™ 480T routing switch for routing to make valuable use of ESRP.

The Layer 2 redundancy features of ESRP offer fast failure recovery (usually four to nine seconds) and provide for multi-homed system design. In some instances, depending on network system design, ESRP can provide better resiliency than using the Spanning Tree Protocol (STP).

For more information on STP, see Chapter 9, "Spanning Tree Protocol (STP)" on page 125.

You can use ESRP instead of STP, but not concurrently with it. You can enable STP on other switches for the VLAN, but the switch configured for ESRP cannot participate in STP for the configured VLAN.

ESRP-Aware Switches

480T routing switches that are not running ESRP, but are connected on a network with other 480T routing switches running ESRP, are ESRP-aware.

When ESRP-aware switches are attached to ESRP-enabled switches, the ESRP-aware switches reliably perform failover and failback scenarios in the prescribed recovery times. It isn't necessary to configure this feature.

When 480T routing switches running ESRP are connected to other types of Layer 2 switches, the failover times for traffic local to the segment may be longer, depending on the application involved and the Forwarding Database (FDB) timer used by the other vendor's Layer 2 switch. As such, you can use ESRP with Layer 2 switches from other vendors, although recovery times vary.

The VLANs associated with the ports connecting an ESRP-aware switch to an ESRP-enabled switch must be configured using an 802.1Q tag on the connecting port or, whenever only a single VLAN is involved, as untagged using the protocol filter `any`.

ESRP Basics

Enterprise Standby Router Protocol (ESRP) is configured on a per-VLAN basis on each switch. A maximum of four switches can participate in providing redundant Layer 3 or Layer 2 services to a single VLAN. A maximum of 64 VLANs can run ESRP simultaneously on a single switch.

The switches exchange keep-alive packets for each VLAN independently. Only one switch can actively provide Layer 3 routing and/or Layer 2 switching for each VLAN. The switch performing the forwarding for a particular VLAN is considered the master for that VLAN. Other participating switches for the VLAN are in standby mode.

For a VLAN with ESRP enabled, each participating switch uses the same MAC address and must be configured with the same IP address or IPX NetID. It is possible for one switch to be master for one or more VLANs while being in standby for others, thus allowing the load to be split across participating switches.

Multiple ESRP VLANs

If multiple ESRP VLANs share a host port, each VLAN must be in an ESRP group.

Mixing Clients and Routers on ESRP VLANs

ESRP should not be enabled on a VLAN that is also expected to exchange routes with other non-ESRP routers (such as routers using RIP or OSPF). ESRP is intended and designed as a Layer 2 or Layer 3 redundancy method for clients with a single default route. ESRP’s fail-over operation may interfere with normal routing protocol communication if an ESRP-enabled VLAN contains other routers not using ESRP.

Ensure that EDP is Enabled

The Enterprise Discovery Protocol (EDP) must be enabled on the ports involved with ESRP in order to function correctly. By default EDP is enabled on all ports. To verify this, use the command:

```
show port <portlist> info
```

To enable EDP on a port, use the command:

```
enable edp ports <portlist>
```

ESRP and Host Attached Ports

Any ESRP VLANs that share ESRP host-attached ports must be in different ESRP groups.

Open Shortest Path First and ESRP

For more information on configuring OSPF, refer to Chapter 13,"RIP and OSPF" on page 223.

If you configure Open Shortest Path First (OSPF) and ESRP, you must manually configure an OSPF router identifier (ID). Be sure that you configure a unique OSPF router ID on each switch running ESRP.

To have two or more switches participate in ESRP, these conditions must be met:

- To make each VLAN redundant, the switches must be able to exchange packets on the same Layer 2 broadcast domain for that VLAN. You can use multiple paths of exchange.

- For a VLAN to be recognized as participating in ESRP, the assigned IP address or the IPX NetID for the separate switches must be identical. Other aspects of the VLAN, including its name, are ignored.
- ESRP must be enabled on the desired VLANs for each switch.

ESRP cannot be enabled on the VLAN default.

- Enterprise Discovery Protocol (EDP) must be enabled on the ports that are members of the ESRP VLANs (the default setting is enabled).

  To verify EDP status, use this command:

  ```
  show ports <portlist> info {detail}
  ```

Determining the ESRP Master

The ESRP master switch (providing Layer 3 routing and/or Layer 2 switching services for a VLAN) is determined by these factors:

- **Active ports**—The switch with the greatest number of active ports takes highest precedence. A load-sharing port group is considered a single port.
- **Tracking information**—In a typical Layer 3 router redundancy configuration (which has the ESRP switches routing to a cloud or routed backbone) you can use the VLAN that links the switch to the routed backbone as part of the criteria for determining the master/slave failover.

  Three types of tracking are supported for determining whether the switch performing the master ESRP function has connectivity to the outside world.

  - VLAN – The number of active ports in a tracked VLAN.
  - IP route – The number of available IP learned routes.
  - Ping – Tracks ICMP ping connectivity to specified devices.

  Other factors being equal, whenever one or more links to the routed backbone fails for the master, ESRP fails-over to the switch that has the most active ports associated with the routed backbone. If there are no active ports associated with the tracked VLAN, ESRP forces the switch to remain in slave state, because no backbone connectivity is available.
- **ESRP priority**—This is a user-defined field. The range of the priority value is 0 to 254; a higher number has higher priority. The

default priority setting is 0. A priority setting of 255 loses the election and remains in standby mode.

- **System MAC address** —The switch with the higher MAC address has priority.

ESRP Tracking

You can use tracking information to monitor various forms of connectivity from the ESRP switch to the outside world. This section describes your ESRP tracking options.

ESRP VLAN Tracking

You can configure ESRP to track connectivity to one or more specified VLANs as criteria for failover. If no active ports remain on the specified VLANs, the switch automatically relinquishes master status and remains in standby mode.

To add or delete a tracked VLAN, use this command:

```
configure vlan <name> [add | delete] track-vlan
<vlan_tracked>
```

ESRP Route Table Tracking

You can configure ESRP to track specified routes in the route table as criteria for failover. If any of the configured routes are not available within the route table, the switch automatically relinquishes master status and remains in standby mode.

To add or delete a tracked route, use this command:

```
configure vlan <name> [add | delete] track-route
<ipaddress/mask_length>
```

ESRP Ping Tracking

You can configure ESRP to track connectivity using a simple ping to any outside responder. The responder may represent the default route of the switch, or any device meaningful to network connectivity of the master ESRP switch.

The switch automatically relinquishes master status and remains in standby mode if a ping keepalive *fails three consecutive times.*

To view the status of tracked devices, use this command:

```
show esrp
```

ESRP Election Algorithms

All switches in the ESRP network must use the same election algorithm, otherwise loss of connectivity, broadcast storms, or other unpredictable behavior may occur.

You configure the switch to use one of five different election algorithms to select the ESRP master. Each algorithm considers the election factors in a different order of precedence, as follows:

- `ports-track-priority-mac`—Active ports, tracking information, ESRP priority, MAC address (Default)
- `track-ports-priority-mac`—Tracking information, active ports, ESRP priority, MAC address
- `priority-ports-track-mac`—ESRP priority, active ports, tracking information, MAC address
- `priority-track-ports-mac`—ESRP priority, tracking information, active ports, MAC address
- `priority-mac-only`—ESRP priority, MAC address

Master Switch Behavior

When a switch is master, it actively provides Layer 3 routing services to other VLANs, and Layer 2 switching between all the ports of that VLAN. Additionally, the switch exchanges ESRP packets with other switches that are in standby mode.

Standby Switch Behavior

When a switch is in standby mode, it exchanges ESRP packets with other switches on that same VLAN. When a switch is in standby, it does not perform Layer 3 routing or Layer 2 switching services for the VLAN.

From a Layer 3 routing protocol perspective (for example, RIP or OSPF), when in standby for the VLAN, the switch marks the router interface associated with the VLAN as down. From a Layer 2 switching perspective, no forwarding occurs between the member ports of the VLAN; this prevents loops and maintains redundancy.

Electing the Master Switch

A new master can be elected in one of these ways:

- A communicated parameter change
- Loss of communication between master and slave(s).

Whenever a parameter that determines the master changes (for example, link loss or priority change), the election of the new master typically occurs within one timer cycle (2 seconds by default).

When a switch in standby mode loses its connection with the master, a new election occurs (using the same precedence order indicated previously).

The new election typically takes place in three times the defined timer cycle (6 seconds by default).

Failover Time

Failover time is largely determined by these factors:

- The ESRP timer setting.
- The routing protocol being used for inter-router connectivity whenever Layer 3 redundancy is used. OSPF failover time is faster than RIP failover time.

The failover time associated with the ESRP protocol depends on the timer setting and the nature of the failure. The default timer setting is 2 seconds; the range is 1 to 255. Default settings usually result in a failover time of 5 to 8.

When routing is configured, the failover of the particular routing protocol (such as RIP V1, RIP V2, or OSPF) is added to the failover time associated with ESRP.

ESRP Options

ESRP options include:

- ESRP Host Attach
- ESRP Domains
- ESRP Groups
- Linking ESRP Switches
- Configuring ESRP and Multinetting
- ESRP and Spanning Tree

ESRP Host Attach

ESRP host attach (HA) is an optional ESRP configuration that allows you to connect active hosts directly to an ESRP master or standby switch.

Normally, the Layer 2 redundancy and loop prevention capabilities of ESRP do not allow packet forwarding from the standby ESRP switch. ESRP HA allows configured ports that do not represent loops to the network to continue Layer 2 operation, independent of their ESRP status.

The ESRP HA option is useful when you are using multi-homed network interface cards (NICs) for server farms, and in conjunction with high-availability server load balancing (SLB) configurations, as shown in Figure 11.1.

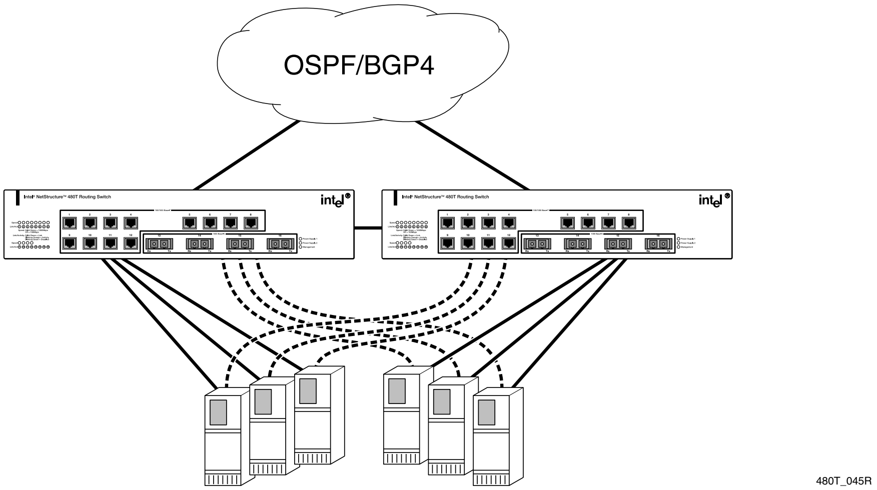

Figure 11.1: ESRP host attach

Other applications allow lower-cost redundant routing configurations, because hosts can be directly attached to the switch involved with ESRP. The ESRP HA feature requires at least one link between the master and standby ESRP switch for carrying traffic and to exchange ESRP hello packets.

ESRP Domains

An ESRP Domain is an optional ESRP configuration that allows you to configure multiple VLANs under the control of a single instance of the ESRP protocol. By grouping multiple VLANs under one ESRP group, the ESRP protocol can scale to provide protection to large numbers of VLANs. All VLANs within an ESRP group simultaneously share the same active and standby router and failover.

ESRP Groups

A switch cannot perform both master and slave functions on the same VLAN for separate instances of ESRP.

The 480T routing switch supports running multiple instances of ESRP within the same VLAN or broadcast domain. This functionality is called an ESRP group. Though other uses exist, the most typical application for multiple ESRP groups is when two or more sets of ESRP switches are providing fast-failover protection within a subnet.

For example, two ESRP switches provide Layer 2 and Layer 3 connectivity and redundancy for the subnet, while another two ESRP switches provide Layer 2 connectivity and redundancy for a portion of the same subnet. Figure 11.2 shows ESRP groups.

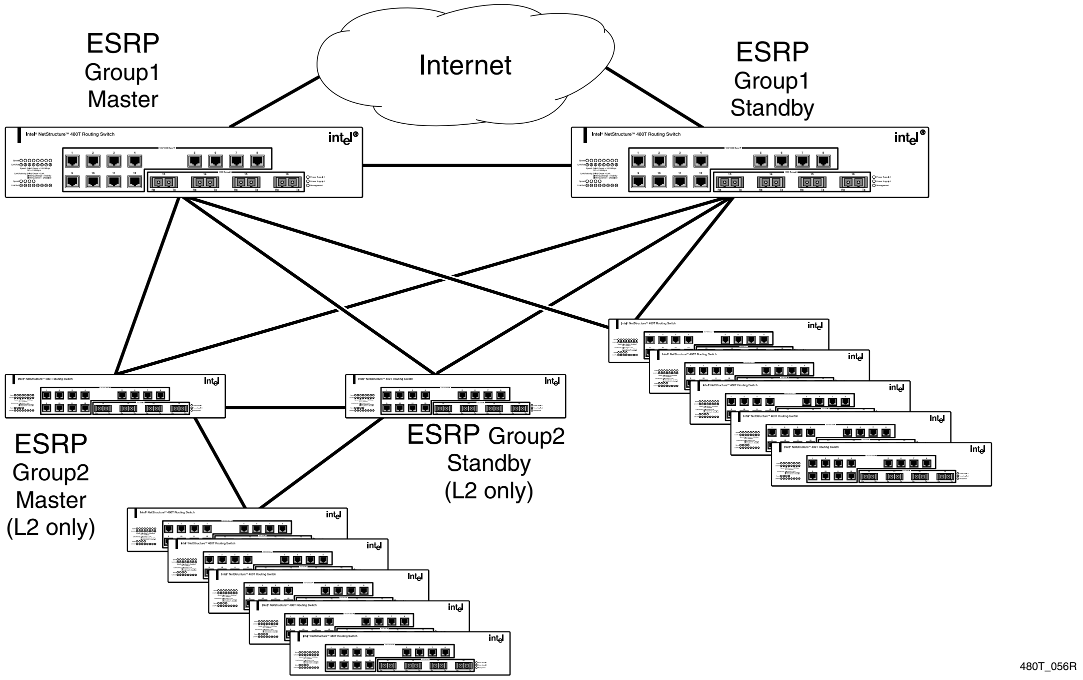

Figure 11.2: ESRP groups

A maximum of four distinct ESRP groups can be supported within the same networked broadcast domain.

To configure the ESRP group membership for a VLAN on a switch, use this command:

```
configure vlan <name> esrp group <group number>
```

The default group number is zero (0).

Linking ESRP Switches

Direct links between ESRP switches are useful under these conditions:

- When the ESRP switches are routing and supporting multiple VLANs (where the master/standby configuration is split so one switch is master for some VLANs and a second switch is master for other VLANs), a direct link provides a more direct path.

  The direct link can contain a unique router-to-router VLAN/subnet. Then the most direct routed path between two VLANs with different master switches is a direct link, instead of forwarding through another set of routers.
- A direct link is a highly reliable method to exchange ESRP hellos, so the possibility of multiple masters for one VLAN is lessened, should all downstream Layer 2 switches fail.
- A direct link is necessary when the ESRP HA option is used. Use the direct link to provide Layer 2 forwarding services through an ESRP standby switch.

Direct links may contain a router-to-router VLAN, along with VLANs running ESRP. When multiple VLANs are used on the direct links, use 802.1Q tagging. The direct links may be aggregated into a load-shared group, if desired.

Configuring ESRP and Multinetting

When configuring ESRP and IP multinetting on the same switch, the parameters that affect the determination of the ESRP master must be configured identically for all the VLANs involved with IP multinetting. For example, the number of links in your configuration, the priority settings, and timer settings must be identical for all affected VLANs.

ESRP and Spanning Tree

A switch running ESRP should not simultaneously participate in Spanning Tree Protocol (STP) for the same VLAN(s). Other switches in the VLAN being protected by ESRP may run STP, in which case, the switch running ESRP forwards the STP BDPUs (Bridge Protocol Data Units), but does not filter them. Therefore, you can combine ESRP and STP on a network and a VLAN, but you must do so on separate devices.

Be careful to maintain ESRP connectivity between ESPR master and standby switches when you design a network that uses ESRP and STP.

ESRP and VLAN Aggregation

Do not configure a sub-VLAN to run ESRP. The system will allow you to enable ESRP on a VLAN and then designate the VLAN as a sub-VLAN, but this is not a supported configuration.

For more information on VLAN aggregation, see Chapter 12, "IP Unicast Routing" on page 189.

You can use ESRP to provide redundant default router protection to VLAN aggregation clients. ESRP is enabled on the super-VLAN only (not the sub-VLANs).

The procedure is to add ports to the super-VLAN that is shared with the sub-VLANs. To do so, configure the super-VLAN with an 802.1Q tag added as tagged with the sub-VLAN ports. This will avoid a protocol conflict. Then enable ESRP on the super-VLAN.

The following example combines ESRP and VLAN aggregation for the super-VLAN ***vsuper*** and two sub-VLANs, ***v1sub*** and ***v2sub***, that have ports 1 and 2 as members, respectively.

1 Create the VLANs and set up the super-VLAN to sub-VLAN relationship:

```
create vlan v1sub
create vlan v2sub
create vlan vsuper
configure vsuper ipaddress 10.1.2.3/24
enable ipforwarding
enable ospf
configure ospf add vsuper
configure v1sub add port 1
configure v2sub add port 2
configure vsuper add subvlan v1sub
configure vsuper add subvlan v2sub
```

2 Turn on ESRP for the VLAN ***vsuper***:

```
configure vsuper tag 1234
configure vsuper add port 1,2 tagged
enable esrp vlan vsuper
```

Use these commands to verify the configuration:

- **show vlan {detail}**—Displays super- and sub-VLAN relationships, IP addresses, and port membership.

- **`show esrp {detail}`**—Verifies ESRP is enabled and operational.

ESRP Commands

Table 11.1 describes the commands used to configure ESRP. Press the Tab key in the command line interface for more command options.

Table 11.1: ESRP Commands

Command	Description
configure esrp port-mode [host | normal] ports <portlist> {dont-count}	Configures the ESRP port mode. A **`normal`** port does not accept or transmit traffic when the local ESRP device is a slave. The host port always switches user traffic, regardless of the ESRP state. The default setting is **`normal`**.
configure vlan <name> add track-diagnostic failover <priority>	Enables the priority of the diagnostic failover.
configure vlan <name> add track-environment failover <priority>	Sets the priority of the environmental failover.
configure vlan <name> add track-ping <ipaddress> frequency <seconds> miss <number>	Configures an ESRP-enabled VLAN to track an external gateway using ping. The switch will not be the ESRP master of the VLAN if the external gateway is not reachable.
configure vlan <name> add track-iproute <ipaddress>/<masklength>	Configures an ESRP-enabled VLAN to track the condition of a route entry in the kernel route table. The switch cannot be the ESRP master if none of the specified routes are reachable.
configure vlan <name> add track-vlan <vlan_tracked>	Configures an ESRP-enabled VLAN to track the condition of another VLAN.

Table 11.1: ESRP Commands (continued)

Command	**Description**
configure vlan <name> delete track-diagnostic	Disables the priority of the diagnostic failover.
configure vlan <name> delete track-environment	Disables the priority of the environmental failover.
configure vlan <name> delete track-ping <ipaddress> frequency <seconds> miss <number>	Configures an ESRP-enabled VLAN to stop tracking an external gateway.
configure vlan <name> delete track-iproute <ipaddress>/<masklength>	Disables route entry tracking for an ESRP-enabled VLAN.
configure vlan <name> delete track-vlan <vlan_tracked>	Removes the tracking of a VLAN by an ESRP-enabled VLAN.
configure vlan <name> esrp esrp-election [ports-track-priority-mac | track-ports-priority-mac | priority-ports-track-mac | priority-track-ports-mac | priority-mac-only]	Configures the election algorithm on the switch. The algorithm must be the same on all switches for a particular VLAN. Specify: • **`ports_track_priority_mac`**—Active ports, tracking information, ESRP priority, MAC address • **`track_ports_priority_mac`**—Tracking information, active ports, ESRP priority, MAC address • **`priority_ports_track_mac`**—ESRP priority, active ports, tracking information, MAC address • **`priority_track_ports_mac`**—ESRP priority, tracking information, active ports, MAC address • **`priority_mac`**—ESRP priority, MAC address The default setting is **`ports_track_priority_mac`**. If no tracking information is configured for a field, the field is ignored.

Table 11.1: ESRP Commands (continued)

Command	Description
configure vlan <name> esrp group <number>	Configures the ESRP group number.
configure vlan <name> esrp priority <value>	Configures the ESRP priority. The range is 0 to 255. The higher number has higher priority. The default setting is `0`. A setting of `255` configures the switch to be in standby state.
configure vlan <name> esrp timer <hello_timer>	Configures the time, in seconds, between ESRP updates. The range is 1 to 255. The default setting is `2`. The timer setting must be configured identically for the VLAN across all participating switches.
configure vlan <name> esrp-group <group number>	Configures the virtual MAC address to be used for the ESRP VLAN. The default group number is `0`.
configure vlan <super_ESRP_VLAN> add domain-member vlan <sub_ESRP_VLAN>	Adds a VLAN to an ESRP domain. ESRP is performed in the domain master VLAN, and not the other domain members. Maximum number of ESRP domain-member VLANs is 3000.
configure vlan <super_ESRP_VLAN> delete domain-member vlan <sub_ESRP_VLAN>	Deletes a VLAN from an ESRP domain.
disable esrp {vlan <name>}	Disables ESRP on a VLAN.
enable esrp vlan <name>	Enables ESRP on a VLAN.
show esrp {detail}	Displays ESRP configuration information.
show esrp vlan <name>	Displays ESRP configuration information for a specific VLAN.

ESRP Examples

This section provides examples of ESRP configurations.

Single VLAN Using Layer 2 and Layer 3 Redundancy

This example, shown in Figure 11.3, uses a number of switches that perform Layer 2 switching for VLAN Sales. The switches are multi-homed to the VLAN Sales switches. The VLAN Sales switches perform Layer 2 switching between the switches shown near the bottom of the diagram, and Layer 3 routing to the outside world.

Each switch is multi-homed using active ports to two VLAN Sales switches (as many as four could be used). ESRP is enabled on each VLAN Sales switch only for the VLAN that connects to the bottom switches.

Each VLAN Sales switch has the VLAN Sales configured using the identical IP address. These switches then connect to the routed enterprise normally, using the desired routing protocol (for example RIP or OSPF).

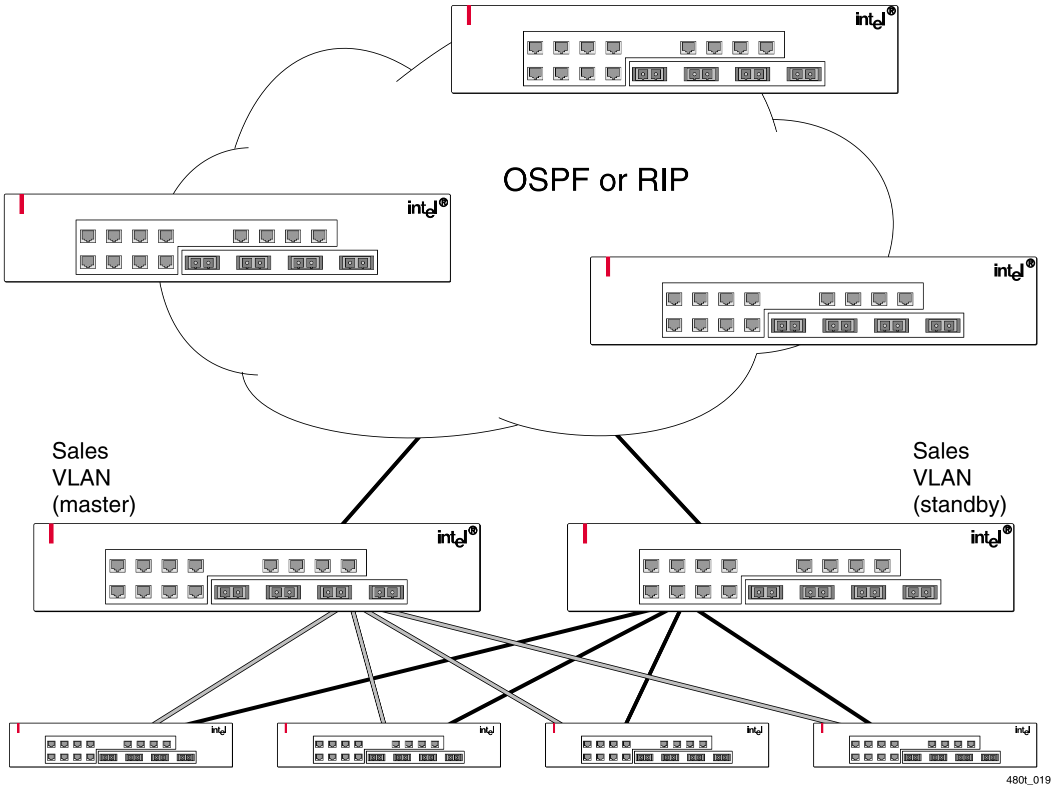

Figure 11.3: ESRP example using Layer 2 and Layer 3 redundancy

The VLAN ***Sales*** master switch, acting as master for VLAN ***Sales***, performs both Layer 2 switching and Layer 3 routing services for VLAN ***Sales***. The switch in standby mode for VLAN ***Sales*** performs neither, thus preventing bridging loops in the VLAN. The switch in standby mode does, however, exchange ESRP packets with the VLAN ***Sales*** master switch.

There are four paths between the VLAN ***Sales*** switches. All the paths are used to send ESRP packets, allowing for four redundant paths for ESRP communication. The switches near the bottom of the diagram, being ESRP-aware, allow traffic within the VLAN to failover quickly, as they will sense when a master/slave transition occurs and flush FDB entries associated with the uplinks to the ESRP-enabled VLAN ***Sales*** switches.

The following commands are used to configure both VLAN ***Sales*** switches. The assumption is that the inter-router backbone is running OSPF, with other routed VLANs already properly configured. Similar commands would be used to configure a switch on a network running RIP. The primary requirement is that the IP address for the VLAN(s) running ESRP must be identical. In this scenario, the master is determined by the programmed MAC address of the switch, because the number of active links for the VLAN and the priority are identical to both switches.

These are the commands used to configure the VLAN ***Sales*** switches:

```
create vlan sales
configure sales add port 1-4
configure sales ipaddr 10.1.2.3/24
enable ipforwarding
enable esrp sales
enable edp ports all
configure ospf add vlan sales
enable ospf
```

Multiple VLANs Using Layer 2 Redundancy

Figure 11.4 illustrates an ESRP configuration that has multiple VLANs using Layer 2 redundancy.

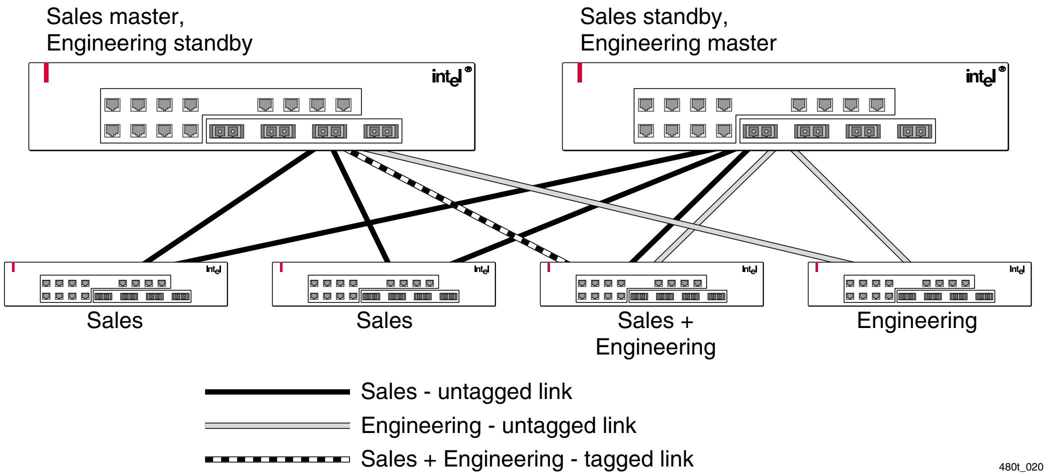

Figure 11.4: ESRP example using Layer 2 redundancy

Figure 11.3 builds on Figure 11.4, but eliminates the requirement of Layer 3 redundancy. It has these features:

- An additional VLAN, ***Engineering***, is added that uses Layer 2 redundancy.
- The VLAN ***Sales*** uses three active links to each upper switch.
- The VLAN ***Engineering*** has two active links to each upper switch.
- The switch labeled ***Sales + Engineering*** carries traffic for both VLANs.
- The link between the ***Sales + Engineering*** switch and the ***Sales master/Engineering standby*** switch uses 802.1Q tagging to carry traffic from both VLANs on one link. The switch counts the link active for each VLAN.
- The ***Sales standby/Engineering master*** switch has a separate physical port for each VLAN connected to the third bottom switch.

In this example, the master and standby switches are configured for ESRP such that the VLAN ***Sales*** normally uses the first switch and the VLAN ***Engineering*** normally uses the second switch. This is accomplished by manipulating the ESRP priority setting for each VLAN for the particular switch.

These are the configuration commands for the first switch (Sales master/Engineering standby):

```
create vlan sales
configure sales tag 10
configure sales add port 1,2
configure sales add port 3 tagged
configure sales ipaddr 10.1.2.3/24
create vlan eng
configure eng tag 20
configure eng add port 4
configure eng add port 3 tagged
configure eng ipaddr 10.4.5.6/24
enable esrp sales
enable esrp eng
enable edp ports all
configure sales esrp priority 5
```

Configuration commands for the second switch (Sales standby/ Engineering master) are as follows:

```
create vlan sales
configure sales add port 1-3
configure sales ipaddr 10.1.2.3/24
create vlan eng
configure eng add port 1,4
configure eng ipaddr 10.4.5.6/24
enable esrp sales
enable esrp eng
configure eng esrp priority 5
```

Displaying ESRP Information

To verify the operational state of an ESRP VLAN and the state of its neighbor, use this command:

```
show esrp
```

To view tracking information about a particular VLAN, including the VLANs tracked by it and a list of the VLANs tracking it, use this command:

```
show vlan
```

ESRP Environment and Diagnostic Tracking

ESRP is capable of tracking hardware status. If a power supply or fan fails, or if the chassis overheats, the priority for the ESRP VLAN will change to the failover settings..

To configure the failover priority for ESRP VLANs, you must first assign a priority to each ESRP VLAN, using this command:

```
configure vlan <vlan name> esrp priority
```

If you set the priority to 255, the ESRP VLAN will remain in standby mode even if the master ESRP VLAN fails. This is a special case.

The range of the priority value is 0 to 254; a higher number has higher priority. The default priority setting is 0.

You typically configure both ESRP VLANs with the same priority.

Next, you must give the priority flag precedence over the active ports count, which has precedence by default, using this command:

```
configure vlan <vlan name> esrp esrp-election
priority-ports-track-mac
```

Since the priority of both VLANs are set equal, ESRP uses the active ports count to determine the master ESRP VLAN.

Then, set the priority of environmental failover using the command:

```
configure vlan <vlan name> add track-environment
failover <priority>
```

Disable the priority of environmental failover, using this command:

```
configure vlan <vlan name> delete track-environment
failover <priority>
```

To enable the priority of the diagnostic failover, use this command:

```
configure vlan <vlan name> add track-diagnostic
failover <priority>
```

To disable the priority of the diagnostic failover, use this command:

```
configure vlan <vlan name> delete track-diagnostic
```

Typically, you set the failover priority lower than the configured priority. Thus, if one of the VLANs experiences a hardware or diagnostics failure, that VLAN becomes the standby VLAN.

12 IP Unicast Routing

This chapter describes how to configure IP routing on the Intel® NetStructure™ 480T routing switch. It assumes that you are already familiar with IP unicast routing. If not, refer to these publications for additional information:

- RFC 1256 — *ICMP Router Discovery Messages*
- RFC 1812 — *Requirements for IP Version 4 Routers*

For IEEE standards information refer to http://standards.ieee.org

Overview of IP Unicast Routing

For more information on routing protocols, refer to "Enterprise Standby Router Protocol (ESRP)" on page 167 and "IP Multicast Routing" on page 275.

The 480T routing switch provides full Layer 3, IP unicast routing. It exchanges routing information with other routers o8n the network using either the Routing Information Protocol (RIP) or the Open Shortest Path First (OSPF) protocol. The 480T routing switch dynamically builds and maintains a routing table and determines the best path for each of its routes.

Each host using the IP unicast routing functionality of the switch must have a unique IP address assigned. In addition, the default gateway assigned to the host must be the IP address of the router interface. RIP and OSPF are described in Chapter 13.

Policy-Based Routing and Route Load-Sharing

Policy-based routing is used to alter the normally calculated next-hop route, which is based on the route table. This same alteration can also load-share across multiple routers. It implies a set of rules or policies that take precedence over information in the route table.

These policies can perform a flow-redirection to different next-hop addresses based on:

- IP source address and mask
- IP destination address and mask
- Layer 4 destination port

In the event that the next-hop address (or addresses) becomes unavailable, the 480T routing switch will route the traffic normally. Several rules may be defined; the precedence of rules is determined by best match of the rule to the packet. If no rule is satisfied, no redirection occurs.

There are two types of commands you can use to set up policy-based routing. One configures the redirection rule(s) and the other configures the next-hop IP address(es):

```
create flow-redirect <flow_rule_name> [tcp | udp]
destination [<ip_address>/<mask> | any] [ip-port
[<L4_port> | any]] source [<ip_address>/<mask> |
any]
```

```
configure flow-redirect <flow_rule_name> [add |
delete] next-hop <ip_address>
```

If multiple next-hop addresses are defined, traffic satisfying the rule is load-shared across the next-hop addresses based on destination IP address.

If next-hop address(es) fail (do not respond to ICMP pings), the switch will resume normal routing.

Using policy-based routing has no impact on switch performance.

To show configuration and status of flow redirection rules, use the command:

```
show flow-redirect [<flow_rule_name | <cr>]
```

Router Interfaces

The routing software and hardware move IP traffic between router interfaces. A router interface is simply a VLAN that has an IP address assigned to it.

As you create VLANs with IP addresses belonging to different IP subnets, you can also choose to route between the VLANs. Both the VLAN switching and IP routing function occur within the 480T routing switch. Each IP address and mask assigned to a VLAN must represent a unique IP subnet.

You cannot configure the same IP address and subnet on different VLANs.

In Figure 12.1, a 480T routing switch is depicted with two VLANs defined, ***Finance*** and ***Personnel***:

- Ports 1 and 3 are assigned to ***Finance***.
- Ports 2 and 4 are assigned to ***Personnel***.
- ***Finance*** belongs to the IP network 192.207.35.0.
- The router interface for ***Finance*** is assigned the IP address 192.207.35.1.
- ***Personnel*** belongs to the IP network 192.207.36.0; its router interface is assigned IP address 192.207.36.1.
- Traffic within each VLAN is switched using the Ethernet MAC addresses.
- Traffic between the two VLANs is routed using the IP addresses.

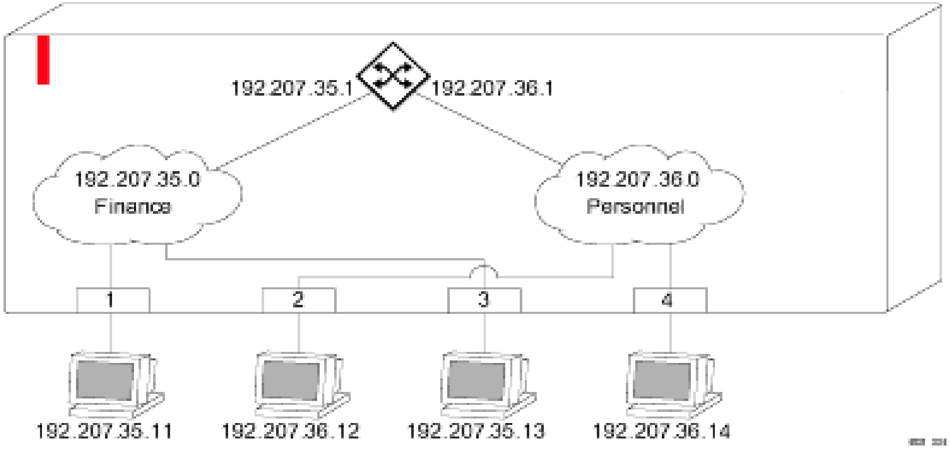

Figure 12.1: Routing between VLANs

Populating the Routing Table

The 480T routing switch maintains an IP routing table for both network routes and host routes. The table is populated from these sources:

If you define a default route, and later delete the VLAN on the subnet associated with it, the now-invalid default route entry remains. You must manually delete the configured default route.

- Dynamically, using routing protocol packets or by ICMP redirects exchanged with other routers
- Statically, using routes entered by the administrator:
  - Default routes, configured by the administrator
  - Locally, using interface addresses assigned to the system
  - By other static routes, as configured by the administrator

Dynamic Routes

Dynamic routes are learned using RIP or OSPF. Routers that use RIP or OSPF exchange information in their routing tables in the form of advertisements. Using dynamic routes, the routing table contains only networks that are reachable.

Dynamic routes are aged out of the table when an update for the network is not received for a period of time, as determined by the routing protocol.

Static Routes

Static routes are manually entered into the routing table. Static routes are used to reach networks not advertised by routers.

Static routes can also be used for security, by controlling which routes you want advertised by the router. You can decide if you want all static routes to be advertised, using one of these commands:

```
[enable | disable] rip export static
```

```
[enable | disable] ospf export static
```

The default setting is enabled. Static routes are never aged out of the routing table.

A static route must be associated with a valid IP subnet. An IP subnet is associated with a single VLAN by its IP address and subnet mask. If the VLAN is subsequently deleted, the static route entries using that subnet must be deleted manually.

Multiple Routes

When there are multiple, conflicting choices of equal-cost routes to a particular destination, the router picks the route with the longest matching network mask. If these are still equal, the router picks the route using these criteria (in this order):

1. Directly attached network interfaces
2. ICMP redirects (refer to Table 12.6 on page 216).
3. Static routes
4. Directly attached network interfaces that are not active

If you define multiple default routes, the route with the lowest metric is used. If there are multiple default routes that have the same lowest metric, the system picks one of the routes.

You can also configure blackhole routes. Traffic to these destinations is silently dropped.

IP Route Sharing

IP route sharing allows multiple equal-cost routes to be used concurrently. You can use IP route sharing with static routes or with OSPF routes. In OSPF, this capability is referred to as equal cost multi-path (ECMP) routing. To use IP route sharing, use this command:

```
enable route sharing
```

Next, configure static routes and/or OSPF as you would normally. You can use as many as five ECMP routes for a given destination.

Route sharing is useful only in instances where you are constrained for bandwidth. This is typically not the case using 480T routing switches. Using route sharing makes router troubleshooting more difficult because of the complexity in predicting the path over which the traffic will travel.

Route Map Support

The switch includes the ability to apply route maps to routes that are being added to the kernel route table. You can configure the route maps based on these origins of the route:

- Direct
- Static
- RIP
- OSPF
- BGP

These route maps match the various characteristics of the route based on the originating protocol and set the accounting indices. Use this command to configure route maps:

```
configure iproute route-map [bgp | direct | e-bgp
| i-bgp | ospf | ospf-extern1 | ospf-extern2 |
ospf-inter | ospf-intra | rip | static] [<route
map> | none]
```

Use this command to view the log:

```
show iproute {priority | vlan <vlan> | permanent |
summary <ipaddress> <netmask> | route-map | origin
[direct | bgp | e-bgp | i-bgp | static | blackhole
| rip | ospf-intra | ospf-inter | ospf-extern1 |
ospf-extern2]} {sorted}
```

You can make dynamic changes to the route map. Direct and static route changes are reflected immediately. RIP changes are reflected within 30 seconds. OSPF and BGP changes depend upon link state.

Route Map Support for OSPF Export

The **enable ospf** command is enhanced to support route maps. The route map is applied on each and every route that is exported to OSPF. It can be used for filtering or for setting the cost, cost type, and tag of the exported route. You can use this feature to make dynamic changes to the route map.

Use these commands to enable OSPF route map export:

```
enable ospf export direct [[cost <metric> [ase-
type-1 | ase-type-2] {tag <number>}] | <route
map>]
```

```
enable ospf export static [[cost <metric> [ase-
type-1 | ase-type-2] {tag <number>} | <route map>]
```

```
enable ospf export rip [[cost <metric> [ase-type-1
| ase-type-2] {tag <number>} | <route map>]
```

```
enable ospf export [bgp | i-bgp | e-bgp] [[cost
<metric> [ase-type-1 | ase-type-2] {tag <number>}
| <route map>]
```

```
enable ospf export vip [[cost <metric> [ase-type-1
| ase-type-2] {tag <number>} | <route map>]
```

BGP and OSPF Route Map Support for Tagging

The 480T routing switch has route map support for BGP and OSPF tagging. This allows you to redistribute OSPF routes from the kernel routing table to BGP, or BGP routes to OSPF.

Use this command to enable tagging:

```
configure route-map <route-map> <sequence number>
[add | delete] match [nlri-list <access-profile> |
as-path [access-profile <access-profile> | <as no>]
| community [access-profile <access-profile> | <as
no>: <number> | number <community> | no-advertise
| no-export | no-export-subconfed] | next-hop <ip
address> | med <number> | origin [igp | egp |
incomplete] | tag <number>]
```

BGP and OSPF Route Map Support for DSB Accounting

The 480T routing switch also offers route map support for BGP and OSPF accounting. This allows you to set the cost and type of the exported routes.

Use this command to enable accounting:

```
configure route-map <route-map> <sequence number>
[add | delete] set [as-path <as no> | community
[[access-profile <access-profile> | <as no>:
<number> | number <community> | no-advertise | no-
export | no-export-subconfed] | remove | [add |
delete] [access-profile <access-profile> | <as no>:
<number> | number <community> | no-advertise | no-
export | no-export-subconfed]] | next-hop <ip
address> | med <number> | local-preference
<number> | weight <number> | origin [igp | egp |
incomplete] | tag <number> | accounting index
<number> value <number> | cost <number> | cost-
type [ase-type-1 | ase-type-2]]
```

Proxy ARP

Proxy Address Resolution Protocol (ARP) was first invented so that ARP-capable devices could respond to ARP request packets on behalf of ARP-incapable devices.

Proxy ARP can also be used to achieve router redundancy and simplify IP client configuration. The switch supports proxy ARP for this type of network configuration.

ARP-Incapable Devices

To configure the switch to respond to ARP requests on behalf of devices that are incapable of doing so, you must configure the IP address and MAC address of the ARP-incapable device using this command:

```
configure iparp add proxy <ipaddress> {<mask>}
<mac_address> {always}
```

Once configured, the system responds to ARP requests on behalf of the device, if these conditions are satisfied:

- The valid IP ARP request is received on a router interface.
- The target IP address matches the IP address configured in the proxy ARP table.
- The proxy ARP table entry indicates that the system should always answer this ARP request, regardless of the ingress VLAN (the **always** parameter must be applied).

Once all the proxy ARP conditions are met, the switch formulates an ARP input rules for the UDP-configured MAC address in the packet.

Proxy ARP Between Subnets

In some networks, it is desirable to configure the IP host with a wider subnet than the actual subnet mask of the segment.

Proxy ARP can be used so that the router answers ARP requests for devices outside of the subnet. As a result, the host communicates as if all devices are local. In reality, communication with devices outside of the subnet are proxied by the router.

For example, an IP host is configured with a class B address of 100.101.102.103 and a mask of 255.255.0.0:

- The switch is configured with the IP address 100.101.102.1 and a mask of 255.255.255.0.
- The switch is also configured with a proxy ARP entry of IP address 100.101.0.0 and mask 255.255.0.0, without the `always` parameter.
- When the IP host tries to communicate with the host at address 100.101.45.67, the IP hosts communicate as if the two hosts are on the same subnet, and sends out an IP ARP request.
- The switch answers on behalf of the device at address 100.101.45.67, using its own MAC address. All subsequent data packets from 100.101.102.103 are sent to the switch, and the switch routes the packets to 100.101.45.67.

Relative Route Priorities

Although these priorities can be changed, do not attempt any manipulation unless you fully understand the possible consequences.

Table 12.1 lists the relative priorities assigned to routes depending on the learned source of the route.

Table 12.1: Relative Route Priorities

Route Origin	Priority
Direct	10
BlackHole	50
Static	1100
ICMP	1200
OSPFIntra	2200
OSPFInter	2300
RIP	2400
OSPFExtern1	3200

Table 12.1: Relative Route Priorities (continued)

Route Origin	Priority
OSPFExtern2	3300
BOOTP	5000

To change the relative route priority, use this command:

```
configure iproute priority [rip | bootp | icmp |
static | ospf-intra | ospf-inter | e-bgp | i-bgp |
ospf-extern1 | ospf-extern2] <priority>
```

IP Multinetting

IP multinetting is used in many legacy IP networks to overlap multiple subnets onto the same physical segment. Due to the resulting constraints in troubleshooting and bandwidth limitations, we recommend that multinetting be used as a transitional tool, and not as a long-term network design strategy.

On the 480T routing switch, each subnet is represented by a different VLAN, and each of those VLANs has its own IP address. All of the VLANs share the same physical ports. The switch routes IP traffic from one subnet to another, all within the same physical ports.

These rules and comments apply when you are configuring IP multinetting:

- Multiple VLANs share the same physical ports; each of the VLANs is configured with an IP address.
- A maximum of four subnets (or VLANs) on multinetted ports is recommended.
- All VLANs used in the multinetting group must share the same port assignment.
- One VLAN is configured to use an IP protocol filter. This is considered the primary VLAN interface for the multinetted group.
- The secondary multinetted VLANs can be exported using the **`export direct`** command.
- The FDB aging timer is automatically set to 3,000 seconds (50 minutes).
- If you are using a UDP or DHCP relay function, only the primary VLAN configured with the IP protocol filter is able to service these requests.
- The VLAN ***default*** should not be used for multinetting.

IP Multinetting Operation

Multinetted VLAN groups must contain identical port assignments.

To use IP multinetting, follow these steps:

1. Select a port where you want IP multinetting to run, for example, port 2.
2. Remove the port from the default VLAN, using this command:

```
configure default delete port 2
```

3. Create a dummy protocol using this command:

```
create protocol mnet
```

4. Create the multinetted subnets using these commands:

```
create vlan net21
create vlan net22
```

5. Assign IP addresses to the net VLANs using these commands:

```
configure net21 ipaddress 123.45.21.1
255.255.255.0
configure net22 ipaddress 192.24.22.1
255.255.255.0
```

6. Assign one of the subnets to the IP protocol using this command:

```
configure net21 protocol ip
```

7. Assign the other subnets to the dummy protocol using this command:

```
configure net22 protocol mnet
```

8. Assign the subnets to a physical port using these commands:

```
configure net21 add port 2
configure net22 add port 2
```

9. Enable IP forwarding on the subnets using this command:

```
enable ipforwarding
```

10. Enable IP multinetting using this command:

```
enable multinetting
```

11. If you are using RIP, disable RIP on the dummy VLANs using this command:

```
configure rip delete net22
```

IP Multinetting Examples

This example configures the switch to have one multinetted segment (port 5) that contains three subnets (192.67.34.0, 192.67.35.0, and 192.67.37.0):

```
configure default delete port 5
create protocol mnet
create vlan net34
create vlan net35
create vlan net37

configure net34 ipaddress 192.67.34.1
configure net35 ipaddress 192.67.35.1
configure net37 ipaddress 192.67.37.1

configure net34 protocol ip
configure net35 protocol mnet
configure net37 protocol mnet

configure net34 add port 5
configure net35 add port 5
configure net37 add port 5

enable ipforwarding
enable multinetting
```

The next example configures the switch to operate with:

- One multinetted segment (port 5) that contains three subnets (192.67.34.0, 192.67.35.0, and 192.67.37.0).
- A second multinetted segment consisting of two subnets (192.67.36.0 and 192.99.45.0). The second multinetted segment spans three ports (port 8, port 9, and port 10).
- RIP enabled on both multinetted segments.

```
configure default delete port 5
create protocol mnet

create vlan net34
create vlan net35
create vlan net37

configure net34 ipaddress 192.67.34.1
configure net35 ipaddress 192.67.35.1
configure net37 ipaddress 192.67.37.1

configure net34 protocol ip
configure net35 protocol mnet
configure net37 protocol mnet

config net34 add port 5
config net35 add port 5
config net37 add port 5

configure default delete port 8,9,10
create vlan net36
create vlan net45
configure net36 ipaddress 192.67.36.1
configure net45 ipaddress 192.99.45.1
configure net36 protocol ip
configure net45 protocol mnet
configure net36 add port 8,9,10
configure net45 add port 8,9,10
configure rip add vlan net34
configure rip add vlan net36
enable ipforwarding
enable multinetting
enable rip
```

Configuring IP Unicast Routing

IGMP and IGMP snooping must be enabled when unicast IP routing or multicast routing is configured (the default setting is enabled for both IGMP and IGMP snooping).

This section describes the commands associated with configuring IP unicast routing on the 480T routing switch. To configure the route:

1. Create and configure two or more VLANs (as described in the previous examples).
2. Assign an IP address to each VLAN that is using routing:

   ```
   configure vlan <name> ipaddress <ipaddress>
   {<mask>}
   ```

   Ensure that each VLAN has a unique IP address.

Default routes are used when the router has no other dynamic or static route to the requested destination.

3. Configure a default route, using this command:

   ```
   configure iproute add default <gateway>
   {<metric>}
   ```

4. Turn on IP routing for one or all VLANs, using this command:

   ```
   enable ipforwarding {vlan <name>}
   ```

5. Turn on RIP or OSPF using one of these commands:

```
enable rip
enable ospf
```

Verifying the IP Unicast Routing Configuration

Use the **`show iproute`** command to display the current configuration of IP unicast routing for the switch and for each VLAN. The **`show iproute`** command displays the currently configured routes and includes how each route was learned.

The **`show iproute`** display has a special flag for routes that are active and in use. These routes are preceded by an asterisk (*) in the route table. If there are multiple routes to the same destination network, the asterisk will indicate which route is preferred.

The Use and M-Use fields in the route table indicate the number of times the software routing module is using the route table entry for packet forwarding decisions.

The Use field indicates a count for unicast routing while the M-Use field indicates a count for multicast routing. If the use count is going up in an unexpected manner, this indicates that the software is making route decisions and may need to be investigated further.

Additional verification commands include:

- **`show iparp`**—Displays the IP ARP table of the system.
- **`show ipfdb`**–Displays the hosts that have been transmitting or receiving packets, and the port and VLAN for each host.
- **`show ipconfig`**–Displays configuration information for one or more VLANs.

VLAN Aggregation

VLAN aggregation is primarily useful to service providers, allowing them to increase the efficiency of IP address space usage. It does this by allowing clients within the same IP subnet to use different broadcast domains using the same default router.

Using VLAN aggregation:

- A superVLAN is defined with the desired IP address, but without any member ports (unless it is running ESRP).

- The subVLANs use the IP address of the superVLAN as the default router address.
- Groups of clients are then assigned to subVLANs that have no IP address, but are members of the superVLAN.
- Clients can be informally allocated any valid IP addresses within the subnet. Optionally, you can prevent communication between subVLANs for isolation purposes so that subVLANs can be quite small, but allow for growth without re-defining subnet boundaries.

Without using VLAN aggregation, each VLAN has a default router address, and you need to use large subnet masks. The result is more unused IP address space.

Multiple secondary IP addresses can be assigned to the superVLAN. These IP addresses are only used to respond to ICMP ping packets to verify connectivity.

Figure 12.2 illustrates VLAN aggregation.

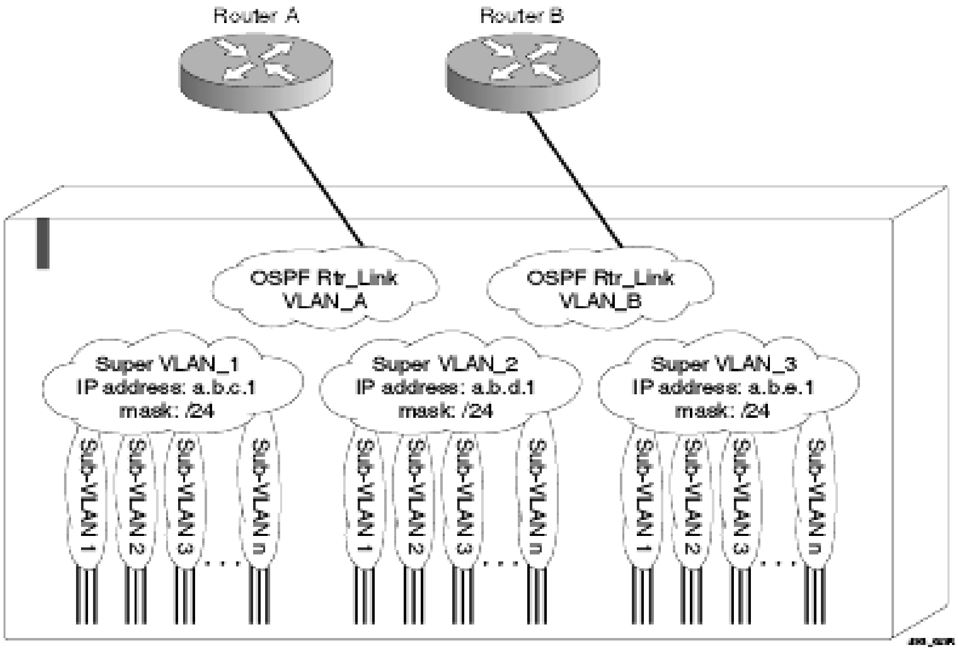

Figure 12.2: VLAN aggregation

In Figure 12.2, all stations are configured to use the address 10.3.2.1 for the default router.

VLAN Aggregation Properties

These properties apply to VLAN aggregation operation:

- All broadcast and unknown traffic remains local to the subVLAN and does not cross the subVLAN boundary.
- All traffic within the subVLAN is switched by the subVLAN, allowing traffic separation between subVLANs (while using the same default router address among them).
- Hosts are located on the subVLAN. Each host can assume any IP address within the address range of the superVLAN router interface. Hosts on the subVLAN are expected to have the same network mask as the superVLAN, and have their default router set to the IP address of the superVLAN.
- All traffic (IP unicast and IP multicast) between subVLANs is routed through the superVLAN. For example, no ICMP redirects are generated for traffic between subVLANs, because the superVLAN is responsible for subVLAN routing.
- Unicast IP traffic across the subVLANs is facilitated by the automatic addition of an ARP entry (similar to a proxy ARP entry) when a subVLAN is added to a superVLAN. You can disable this feature for security purposes.
- IP multicast traffic between subVLANs is routed when an IP multicast routing protocol is enabled on the superVLAN.

VLAN Aggregation Limitations

These limitations apply to VLAN aggregation:

- No additional routers may be located in a subVLAN.
- A subVLAN cannot be a superVLAN, and vice-versa.
- subVLANs are not assigned IP addresses.
- Typically, a superVLAN has no ports associated with it, except in the case of running ESRP.
- If a client is moved from one subVLAN to another, clear the IP ARP cache at the client and the switch to resume communication.

SubVLAN Address Range Checking

The use of static ARP entries associated with superVLANs or sub-VLANs is not supported.

Sub-VLAN address ranges can be configured on each subVLAN to prohibit the entry of IP addresses from hosts outside of the configured range.

To configure a subVLAN range use this command:

```
configure vlan <vlan_name> subvlan-address-range
<ipaddress-ipaddress>
```

To remove a subVLAN address range use this command:

```
configure vlan <vlan_name> subvlan-address-range
0.0.0.0 - 0.0.0.0
```

To view a subVLAN range use this command:

```
show vlan [vlan_name]
```

There is no error checking to prevent the configuration of overlapping subVLAN address ranges between multiple subVLANs. Doing so may result in unpredictable behavior of ARP within the superVLAN and associated subVLANs.

Isolation Option for Communication Between subVLANs

The isolation option works for normal, dynamic, ARP-based client communication.

To facilitate communication between subVLANs, by default, an entry is made in the IP ARP table of the superVLAN that performs a proxy ARP function. This allows clients on one subVLAN to communicate with clients on another subVLAN. In certain circumstances, intra-subVLAN communication may not be desired for isolation reasons.

To prevent normal communication between subVLANs, disable the automatic addition of the IP ARP entries on the superVLAN, using the command:

```
disable subvlan-proxy-arp vlan <supervlan name>
```

VLAN Aggregation Commands

Table 12.2 describes VLAN aggregation commands. For more command options press the Tab key in the command line interface.

Table 12.2: VLAN Aggregation Commands

Command	Description
configure vlan <supervlan name> add secondary-ip <ipaddress> {<mask>}	Adds a secondary IP address to the superVLAN for responding to ICMP ping requests.
configure vlan <supervlan name> add subvlan <subvlan name>	Adds a subVLAN to a superVLAN.
configure vlan <supervlan name> delete secondary-ip <ipaddress> {<mask>}	Deletes a secondary IP address to the superVLAN for responding to ICMP ping requests.
configure vlan <supervlan name> delete subvlan <subvlan name>	Deletes a subVLAN from a superVLAN.
disable subvlan-proxy-arp vlan [<supervlan name> | all]	Disables subVLAN entries in the proxy ARP table.
enable subvlan-proxy-arp vlan [<supervlan name> | all]	Enables the automatic entry of subVLAN information in the proxy ARP table.

VLAN Aggregation Example

The following example illustrates how to configure VLAN aggregation. The VLAN ***vsuper*** is created as a superVLAN, and subVLANs ***vsub1***, ***vsub2***, and ***vsub3*** are added to it.

1. Create and assign an IP address to a VLAN designated as the superVLAN. This VLAN should have no member ports. Be sure to enable IP forwarding, and any desired routing protocol, on the switch:

```
create vlan vsuper
configure vsuper ipaddress 192.201.3.1/24
enable ipforwarding
enable ospf
configure ospf add vsuper
```

2. Create and add ports to the subVLANs:

```
create vlan vsub1
configure vsub1 add port 8-10
create vlan vsub2
configure vsub2 add port 11-13
create vlan vsub3
configure vsub3 add port 15-16
```

3. Configure the superVLAN by adding the subVLANs:

```
configure vsuper add subvlan vsub1
configure vsuper add subvlan vsub2
configure vsuper add subvlan vsub3
```

4. Disable communication among subVLANs (optional):

```
disable subvlan-proxy-arp <superVLAN name>
```

Verifying the VLAN Aggregation Configuration

Use these commands to verify proper VLAN aggregation configuration:

- `show vlan`—Indicates the membership of subVLANs in a superVLAN.
- `show iparp`—Indicates an ARP entry that contains subVLAN information. Communication with a client on a subVLAN must occur before you can make an entry in the ARP table.

Configuring DHCP/BOOTP Relay

Once IP unicast routing is configured, you can configure the switch to forward Dynamic Host Configuration Protocol (DHCP) or BOOTP requests coming from clients on subnets being serviced by the switch and going to hosts on different subnets.

You can use this feature in various applications, including DHCP services between Windows§ NT§ servers and clients running Windows 95. To configure the relay function:

- Configure VLANs and IP unicast routing.
- Enable the DHCP or BOOTP relay function, using this command:

```
enable bootprelay
```

- Configure the addresses, where you want to direct DHCP or BOOTP requests, using this command:

  ```
  configure bootprelay add <ipaddress>
  ```

To delete an entry, use this command:

```
configure bootprelay delete {<ipaddress> | all}
```

Verifying the DHCP/BOOTP Relay Configuration

To verify the DHCP/BOOTP relay configuration, use this command:

```
show ipconfig
```

This command displays the configuration of the BOOTP relay service, and the addresses that are currently configured.

UDP Forwarding

UDP forwarding is a flexible and generalized routing utility for handling the directed forwarding of broadcast UDP packets. UDP forwarding allows applications, such as multiple DHCP relay services from differing sets of VLANs, to be directed to different DHCP servers. These rules apply to UDP broadcast packets handled by this feature:

- If the UDP profile includes BOOTP or DHCP, it is handled according to guidelines in RFC 1542.
- If the UDP profile includes other traffic types, these packets have the IP destination address modified as configured, and changes are made to the IP and UDP checksums and decrements to the TTL field, as appropriate.

If the UDP forwarding is used for BOOTP or DHCP forwarding purposes, do not configure or use the existing **`bootprelay`** function. However, if the previous **`bootprelay`** functions are applicable, you may continue to use them.

Configuring UDP Forwarding

To configure UDP forwarding, the you must first create a UDP-forward destination profile. The profile describes the types of UDP packets (by port number) that are used, and where they are to be forwarded. You must give the profile a unique name, in the same manner as a VLAN, protocol filter, or Spanning Tree Domain.

Next, configure a VLAN to use the UDP-forwarding profile. As a result, all incoming traffic from the VLAN that matches the UDP profile is handled as specified in the UDP-forwarding profile.

You can define a maximum of ten UDP-forwarding profiles. Each named profile may contain a maximum of eight rules defining the UDP port, and destination IP address or VLAN. A VLAN can use a single UDP-forwarding profile. UDP packets directed toward a VLAN use an all-ones broadcast on that VLAN.

UDP-Forwarding Example

In this example, the VLAN ***Marketing*** and the VLAN ***Operations*** are pointed toward a specific backbone DHCP server (with IP address 10.1.1.1) and a backup server (with IP address 10.1.1.2). Additionally, the VLAN ***LabUser*** is configured to use any responding DHCP server on a separate VLAN called ***LabSvrs***.

The commands for this configuration are:

```
create udp-profile backbonedhcp
create udp-profile labdhcp
configure backbonedhcp add 67 ipaddress 10.1.1.1
configure backbonedhcp add 67 ipaddress 10.1.1.2
configure labdhcp add 67 vlan labsvrs
configure marketing udp-profile backbonedhcp
configure operations udp-profile backbonedhcp
configure labuser udp-profile labdhcp
```

ICMP Packet Processing

As ICMP packets are routed or generated, you have several options for controlling distribution:

Access lists are described in Chapter , "Access Policies" on page 309.

- For ICMP packets typically generated or observed as part of the routing function, you can assert control on a per-type, per-VLAN basis.

- You can alter the default settings for security reasons, by restricting the success of tools that could be used to find information on an important application, host, or topology.
- For ICMP packets that are typically routed, you can apply access lists to restrict forwarding behavior.

The controls include the disabling of transmitting ICMP messages associated with:

- unreachables
- port-unreachables
- time-exceeded
- parameter-problems
- redirects
- time-stamp
- address-mask requests.

UDP-Forwarding Commands

Table 12.3 describes the commands used to configure UDP-forwarding. For more command options, press the Tab key in the command line interface.

Table 12.3: UDP-Forwarding Commands

Command	Description
configure udp-profile <profile_name> add <udp_port> [vlan <name> | ipaddress <dest_ipaddress>]	Adds a forwarding entry to the specified User Datagram Protocol(UDP)-forwarding profile name. All broadcast packets sent to **`udp_port`** are forwarded to either the destination IP address (unicast or subnet directed broadcast) or to the specified VLAN as an all-ones broadcast.
configure udp-profile <profile_name> delete <udp_port> [vlan <name> | ipaddress <dest_ipaddress>]	Deletes a forwarding entry from the specified **`udp-profile`** name.

Table 12.3: UDP-Forwarding Commands (continued)

Command	**Description**
configure vlan <name> udp-profile <profile_name>	Assigns a UDP-forwarding profile to the source VLAN. Once the UDP profile is associated with the VLAN, the Intel® NetStructure™ 480T routing switch picks up any broadcast UDP packets that match the user-configured UDP port number, and forwards those packets to the user-defined destination. If the UDP port is the DHCP/BOOTP port number, appropriate DHCP/BOOTP proxy functions are used.
create udp-profile <profile_name>	Creates a UDP-forwarding profile. You must use a unique name for the UDP-forwarding profile.
delete udp-profile <profile_name>	Deletes a UDP-forwarding profile.
show udp-profile {<profile_name>}	Displays the profile names, input rules of the UDP port, destination IP address, or VLAN and the source VLANs where the profile is applied.
unconfigure udp-profile vlan [<name> | all]	Removes the UDP-forwarding profile configuration for one or all VLANs.

IP Commands

Table 12.4 describes the commands used to configure basic IP settings. For more command options, press the Tab key in the command line interface.

Table 12.4: Basic IP Commands

Command	Description
clear iparp {<ipaddress> | vlan <name>}	Removes dynamic entries in the IP ARP table. Permanent IP ARP entries are not affected.
clear ipfdb {<ipaddress> | vlan <name> }	Removes the dynamic entries in the IP forwarding database. If no options are specified, all dynamic IP FDB entries are removed.
configure bootprelay add <ipaddress>	Adds the IP destination address to forward BOOTP packets.
configure bootprelay delete [<ipaddress> | all]	Removes one or all IP destination addresses for forwarding BOOTP packets.
configure iparp add <ipaddress> <mac_address>	Adds a permanent entry to the ARP table. Specify the IP address and MAC address of the entry.
configure iparp add proxy <ipaddress> {<mask>} {<mac_address>} {always}	Configures proxy ARP entries. When **`mask`** is not specified, an address with the mask `255.255.255.255` is assumed. When **`mac_address`** is not specified, the MAC address of the Intel® NetStructure™ 480T routing switch is used in the ARP response. When **`always`** is specified, the switch answers ARP requests without filtering requests that belong to the same subnet of the receiving router interface.
configure iparp delete <ipaddress>	Deletes an entry from the ARP table. Specify the IP address of the entry.
configure iparp delete proxy [<ipaddress> {<mask>} | all]	Deletes one or all proxy ARP entries.

Table 12.4: Basic IP Commands (continued)

Command	Description
configure iparp timeout <minutes>	Configures the IP ARP timeout period. The default setting is 20 minutes. A setting of 0 disables ARP aging. The maximum aging time is 32 minutes.
configure tcp-sync-rate <number_sync_per_sec>	Configures a limit for the switch to process TCP connection requests. If the connection request rate is higher than the specified rate, or the total number of outstanding connection requests exceeds the system limit, the system ages out incomplete connection requests at a faster rate. The range is 5 to 200,000. The default setting is 25 connection requests per second.
disable bootp vlan [<name> | all]	Disables the generation and processing of BOOTP packets.
disable bootprelay	Disables the forwarding of BOOTP requests.
disable ipforwarding {vlan <name>}	Disables routing for one or all VLANs.
disable ipforwarding broadcast {vlan <name>}	Disables routing of broadcasts to other networks.
disable loopback-mode vlan [<name>]	Disables loopback mode on an interface.
disable multinetting	Disables IP multinetting on the system.
enable bootp vlan [<name> | all]	Enables the generation and processing of BOOTP packets on a VLAN to obtain an IP address for the VLAN from a BOOTP server. The default setting is enabled for all VLANs.

Table 12.4: Basic IP Commands (continued)

Command	Description
enable bootprelay	Enables the forwarding of BOOTP and Dynamic Host Configuration Protocol (DHCP) requests.
enable ipforwarding {vlan <name>}	Enables IP routing for one or all VLANs. If no argument is provided, enables routing for all VLANs that are configured with an IP address. The default setting for **`ipforwarding`** is disabled.
enable ipforwarding broadcast {vlan <name>}	Enables forwarding IP broadcast traffic for one or all VLANs. If no argument is provided, enables broadcast forwarding for all VLANs. To enable, **`ipforwarding`** must be enabled on the VLAN. The default setting is disabled.
enable loopback-mode vlan [<name>]	Enables a loopback mode on an interface. If loopback is enabled, the router interface remains in the UP state, even if no ports are defined in the VLAN. As a result, the subnet is always advertised as one of the available routes.
enable multinetting	Enables IP multinetting on the system.

Table 12.5 describes the commands used to configure the IP route table. For more command options, press the Tab key in the command line interface.

Table 12.5: Route Table Configuration Commands

Command	Description
configure iproute add <ipaddress>/<mask> <gateway> <metric>	Adds a static address to the routing table. Use a value of 255.255.255.255 for **`mask`** to indicate a host entry.

Table 12.5: Route Table Configuration Commands (continued)

Command	Description
configure iproute add blackhole <ipaddress> <mask>	Adds a blackhole address to the routing table. All traffic destined for the configured IP address is dropped, and no Internet Control Message Protocol (ICMP) message is generated.
configure iproute add default <gateway> {<metric>}	Adds a default gateway to the routing table. A default gateway must be located on a configured IP interface. If no metric is specified, the default metric of 1 is used. Use the unicast-only or multicast-only options to specify a particular traffic type. If not specified, both unicast and multicast traffic use the default route.
configure iproute delete <ipaddress> <mask> <gateway>	Deletes a static address from the routing table.
configure iproute delete blackhole <ipaddress> <mask>	Deletes a **blackhole** address from the routing table.
configure iproute delete default <gateway>	Deletes a default gateway from the routing table.
configure iproute priority [rip | bootp | icmp | static | ospf-intra | ospf-inter | e-bgp | i-bgp | ospf-extern1 | ospf-extern2] <priority>	Changes the priority for all routes coming from a particular route origin.
disable iproute sharing	Disables load sharing for multiple routes.
enable iproute sharing	Enables load sharing if multiple routes to the same destination are available. Only paths with the same lowest cost are shared. The default setting is disabled.
rtlookup [<ipaddress> | <hostname>]	Performs a look-up in the route table to determine the best route to reach an IP address.

Table 12.6 describes the commands used to configure IP options and the ICMP protocol. For more command options, press the Tab key in the command line interface.

Table 12.6: ICMP Configuration Commands

Command	Description
configure irdp [multicast | broadcast]	Configures the destination address of the router advertisement messages. The default setting is `multicast`.
configure irdp <mininterval> <maxinterval> <lifetime> <preference>	Configures the router advertisement message timers, in seconds. Specify: `mininterval`—The minimum amount of time between router advertisements. The default setting is `450`. `maxinterval`—The maximum time between router advertisements. The default setting is `600`. `lifetime`—The default setting is `1,800`. `preference`**—**The preference level of the router. An ICMP Router Discover Protocol (IRDP) client always uses the router with the highest preference level. Change this setting to encourage or discourage the use of this router. The default setting is `0`.
disable icmp [parameter-problem | address-mask | port-unreachables | redirects | time-exceeded | timestamp | unreachables | userdirects] {vlan <name>}	Disables ICMP messages for the packet type specified.
disable ip-option loose-source-route	Disables the loose source route IP option.
disable ip-option record-route	Disables the record-route IP option.
disable ip-option record-timestamp	Disables the record-timestamp IP option.
disable ip-option strict-source-route	Disables the strict-source-route IP option.
disable ip-option use-router-alert	Disables the use router alert IP option.

Table 12.6: ICMP Configuration Commands (continued)

Command	Description
enable icmp address-mask vlan [<name> | all]	Enables an ICMP address-mask reply (type 18, code 0) when an ICMP address mask request is received.The default setting is enabled. If a VLAN is not specified, the command applies to all IP interfaces.
enable icmp parameter-problem vlan [<name> | all]	Enables an ICMP parameter problem message packet (type 12) when the switch cannot properly process the IP header or IP option information.
enable icmp port-unreachables vlan [<name> | all]	Enables ICMP port unreachable messages (type 3, code 3) when a TPC or UDP request is made to the switch, and no application is waiting for the request, or when access policy denies the request. The default setting is enabled. If a VLAN is not specified, the command applies to all IP interfaces.
enable icmp redirects vlan [<name> | all]	Enables an ICMP redirect message (type 5) when a packet must be forwarded out on the ingress port. The default setting is enabled. If a VLAN is not specified, the command applies to all IP interfaces.
enable icmp time-exceeded vlan [<name> | all]	Enables an ICMP time exceeded message (type 11) when the TTL field expires during forwarding. IP multicast packets do not trigger ICMP time exceeded messages. The default setting is enabled. If a VLAN is not specified, the command applies to all IP interfaces.
enable icmp timestamp vlan [<name> | all]	Enables an ICMP timestamp response (type 14, code 0) when an ICMP timestamp request is received. The default setting is enabled. If a VLAN is not specified, the command applies to all IP interfaces.

Table 12.6: ICMP Configuration Commands (continued)

Command	Description
enable icmp unreachables vlan [<name> | all]	Enables ICMP network-unreachable messages (type 3, code 0), and host unreachable messages (type 3, code 1) when a packet cannot be forwarded to the destination because of an unreachable route or host. The default setting is enabled. If a VLAN is not specified, the command applies to all IP interfaces.
enable icmp useredirects	Enables the modification of route table information when an ICMP redirect message is received. This option applies to the switch when it is *not configured for routing*. The default setting is disabled.
enable ip-option loose-source-route	Enables the loose source route IP option.
enable ip-option record-route	Enables the record route IP option.
enable ip-option record-timestamp	Enables the record timestamp IP option.
enable ip-option strict-source-route	Enables the strict source route IP option.
enable ip-option use-router-alert	Enables the Intel® NetStructure™ 480T routing switch to generate the router alert IP option with routing protocol packets.
enable irdp {vlan <name>}	Enables ICMP router advertisement messages on one or all VLANs. The default setting is enabled.
unconfigure icmp	Resets all ICMP settings to the default values.
unconfigure irdp	Resets all router-advertisement settings to the default values.

Routing Configuration Example

Figure 12.3 illustrates a 480T routing switch with three VLANs defined as:

- ***Financeaddress 192.207.35.1***
  - protocol sensitive VLAN using the IP protocol
  - Ports 1 and 3 are assigned
  - IP address 192.207.35.1.
- ***Personnel***
  - Protocol-sensitive VLAN using the IP protocol
  - Ports 2 and 4 are assigned
  - IP address 192.207.36.1
- ***MyCompany***
  - Port-based VLAN
  - All ports are assigned

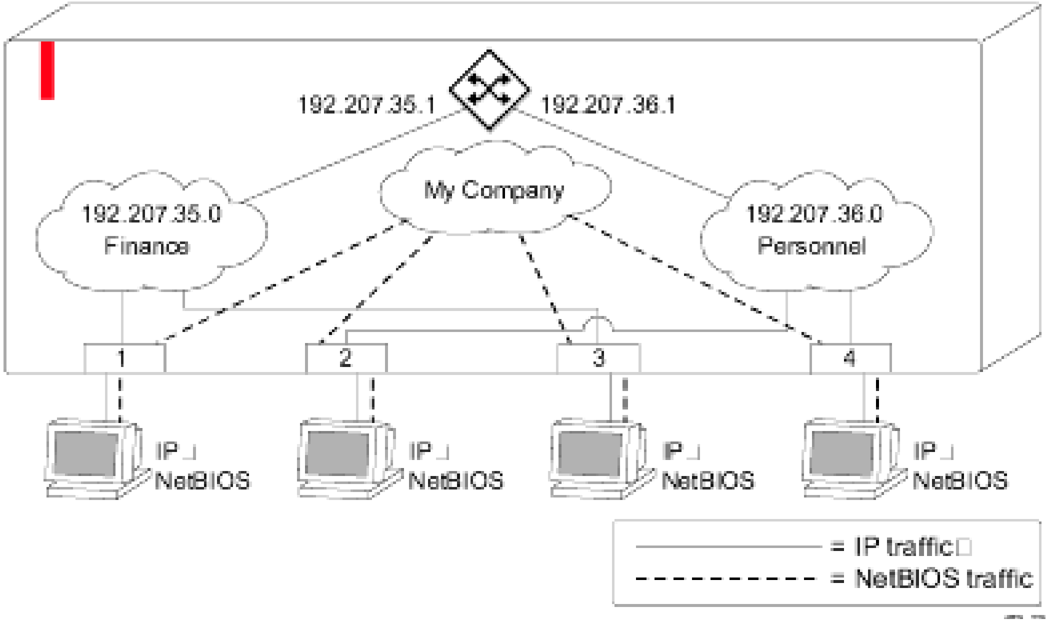

Figure 12.3: Unicast routing configuration example

The stations connected to ports 1 through 4 generate a combination of IP traffic and NetBIOS traffic. The IP traffic is filtered by the protocol-sensitive VLANs. All other traffic is directed to the VLAN ***MyCompany***.

In this configuration, all IP traffic from stations connected to ports 1 and 3 have access to the router using the VLAN ***Finance***. Ports 2 and 4 reach the router through the VLAN ***Personnel***. All other traffic (NetBIOS) is part of the VLAN ***MyCompany***.

The example in Figure 12.3 is configured using these commands:

```
create vlan Finance
create vlan Personnel
create vlan MyCompany

configure Finance protocol ip
configure Personnel protocol ip

configure Finance add port 1,3
configure Personnel add port 2,4
configure MyCompany add port all

configure Finance ipaddress 192.207.35.1
configure Personnel ipaddress 192.207.36.1

configure rip add vlan Finance
configure rip add vlan Personnel

enable ipforwarding
enable rip
```

Displaying Router Settings

To display settings for various IP routing components, use the commands listed in Table 12.7. For more command options, press the Tab key in the command line interface.

Table 12.7: Router Show Commands

Command	Description
show iparp proxy {<ipaddress> {<mask>}}	Displays the proxy Address Resolultion Protocol (ARP) table.

Table 12.7: Router Show Commands (continued)

Command	Description
show iparp {<ipaddress | vlan <name> | permanent}	Displays the IP ARP table. You can filter the display by IP address, VLAN, or permanent entries.
show ipconfig {vlan <name> | detail}	Displays configuration information for one or all VLANs.
show ipfdb {<ipaddress>/<netmask> | vlan <name>}	Displays the contents of the IP FDB table. If no option is specified, all IP FDB entries are displayed.
show iproute {priority | vlan <name> | permanent | <ipaddress> <mask> | origin [direct | static | blackhole | rip | ospf-intra | ospf-inter | ospf-extern1 bgp | e-bgp | i-bgp | ospf-extern2]} {sorted}	Displays the contents of the IP routing table or the route origin priority.
show ipstats {vlan <name>}	Displays IP statistics for the microprocessor.

Resetting and Disabling Router Settings

To return router settings to their defaults and disable routing functions, use the commands listed in Table 12.8. For more command options, press the Tab key in the command line interface.

Table 12.8: Router Reset and Disable Commands

Command	Description
clear iparp {<ipaddress> | vlan <name>}	Removes dynamic entries in the IP ARP table. Permanent IP ARP entries are not affected.
clear ipfdb {<ipaddress> | vlan <name>]	Removes the dynamic entries in the IP forwarding database. If no options are specified, all IP FDB entries are removed.

Table 12.8: Router Reset and Disable Commands (continued)

Command	Description
disable bootp vlan [<name> | all]	Disables the generation and processing of BOOTP packets.
disable bootprelay	Disables the forwarding of BOOTP requests.
disable icmp <packet-type> vlan [<name>]	Disables ICMP parameter-problem messages. If a VLAN is not specified, the command applies to all IP interfaces. Packet types are: **`parameter-problem, address-mask, port-unreachable, redirect, time-exceeded, timestamp, unreachable, userdirect`**.
disable ipforwarding broadcast {vlan <name>}	Disables routing of broadcasts to other networks.
disable ipforwarding {vlan <name>}	Disables routing for one or all VLANs.
disable irdp {vlan <name>}	Disables router advertisement messages on one or all VLANs.
unconfigure icmp	Resets all ICMP settings to the default values.
unconfigure irdp	Resets all router advertisement settings to the default values.

13 RIP and OSPF

This chapter describes the interior routing protocols available on the Intel® NetStructure™ 480T routing switch, RIP and OSPF. It assumes that you are already familiar with IP unicast routing. If not, refer to these publications:

- RFC 1058 — *Routing Information Protocol (RIP)*
- RFC 1723 — *RIP Version 2*
- RFC 2178 — *OSPF Version 2*

Overview

Both RIP and OSPF can be enabled on a single VLAN.

The Intel NetStructure 480T routing switch supports the use of two interior gateway protocols (IGPs): the Routing Information Protocol (RIP) and the Open Shortest Path First (OSPF) protocol for IP unicast routing.

RIP is a distance-vector protocol, based on the Bellman-Ford (or distance-vector) algorithm. The distance-vector algorithm has been in use for many years, and is widely deployed and understood.

OSPF is a link-state protocol, based on the Dijkstra link-state algorithm. OSPF is a newer IGP, and solves a number of problems associated with using RIP on today’s complex networks.

Both RIP and OSPF can be enabled on a single VLAN.

Distinguishing RIP and OSPF

The distinction between RIP and OSPF lies in the fundamental differences between distance-vector protocols and link-state protocols. Using a distance-vector protocol, each router creates a unique routing table from summarized information obtained from neighboring routers. Using a link-state protocol, every router maintains an identical routing table created from information obtained from all routers in the autonomous system. Each router builds a shortest path tree, using itself as the root. The link-state protocol ensures that updates sent to neighboring routers are acknowledged by the neighbors, verifying that all routers have a consistent network map.

The biggest advantage of using RIP is that it is relatively simple to understand and implement, and it is the de facto routing standard for many years.

RIP has a number of limitations that can cause problems in large networks, including:

- A limit of 15 hops between the source and destination networks
- A large amount of bandwidth taken up by periodic broadcasts of the entire routing table
- Slow convergence
- Routing decisions based on hop count; no concept of link costs or delay
- Flat networks; no concept of areas or boundaries

OSPF offers many advantages over RIP, including:

- No limitation on hop count
- Route updates multicast only when changes occur
- Faster convergence
- Support for load balancing to multiple routers based on the actual cost of the link
- Support for hierarchical topologies where the network is divided into areas

The details of RIP and OSPF are explained later in this chapter.

Overview of RIP

RIP is an IGP first used in computer routing in the Advanced Research Projects Agency Network (ARPAnet) as early as 1969. It is primarily intended for use in homogeneous networks of moderate size.

To determine the best path to a distant network, a router using RIP always selects the path that has the least number of hops. Each router that data must traverse is considered to be one hop.

Routing Table

The routing table in a router using RIP contains an entry for every known destination network. Each routing table entry contains:

- IP address of the destination network
- Metric (hop count) to the destination network
- IP address of the next router
- Timer that tracks the amount of time since the entry was last updated

The router exchanges an update message with each neighbor every 30 seconds (default value), or if there is a change to the overall routed topology (also called triggered updates). If a router does not receive an update message from its neighbor within the route timeout period (180 seconds by default), the router assumes the connection between it and its neighbor is no longer available.

Split Horizon

Split horizon is a scheme for avoiding problems caused by including routes in updates sent to the router from which the route was learned. Split horizon omits routes learned from a neighbor in updates sent to that neighbor.

Poison Reverse

Like split horizon, poison reverse is a scheme for eliminating the possibility of loops in the routed topology. In this case, a router advertises a route over the same interface that supplied the route, but the route uses a hop count of 16, defining it as unreachable.

Triggered Updates

Triggered updates occur whenever a router changes the metric for a route, and it is required to send an update message immediately, even if it is not yet time for a regular update message to be sent. This will generally result in faster convergence, but may also result in more RIP-related traffic.

Route Advertisement of VLANs

VLANs that are configured with an IP address, but are configured to not route IP or are not configured to run RIP, do not have their subnets advertised by RIP. Only those VLANs that are configured with an IP address and are configured to route IP and run RIP have their subnets advertised.

RIP Version 1 Compared to RIP Version 2

When using RIP with supernetting/Classless Inter-Domain Routing (CIDR), use RIPv2 only. Turn RIP route aggregation off.

RIP version 2 expands the functionality of RIP version 1 to include:

- Variable-Length Subnet Masks (VLSMs)
- Support for next-hop addresses, which allows for optimization of routes in certain environments.
- Multicasting; RIP version 2 packets can be multicast instead of being broadcast, reducing the load on hosts that do not support routing protocols.

Overview of OSPF

OSPF is a link-state protocol that distributes routing information between routers belonging to a single IP domain, also known as an autonomous system (AS). In a link-state routing protocol, each router maintains a database describing the topology of the autonomous system. Each participating router has an identical database maintained from the perspective of that router.

From the link-state database (LSDB), each router constructs a tree of shortest paths, using itself as the root. The shortest path tree provides the route to each destination in the autonomous system. When several equal-cost routes to a destination exist, traffic can be distributed among them. The cost of a route is described by a single metric.

Link-State Database

Upon initialization, each router transmits a link-state advertisement (LSA) on each of its interfaces. LSAs are collected by each router and entered into the LSDB of each router. OSPF uses flooding to distribute LSAs between routers. Any change in routing information is sent to all of the routers in the network. All routers within an area have the exact same LSDB. Table 13.1 describes LSA type numbers.

Table 13.1: LSA Type Numbers

Type Number	Description
1	Router LSA
2	Network LSA
3	Summary LSA
4	AS summary LSA
5	AS external LSA
7	Not-so-stubby-area (NSSA) external LSA

Areas

OSPF allows parts of a network to be grouped together into areas. The topology within an area is hidden from the rest of the autonomous system. Hiding this information enables a significant reduction in LSA traffic, and reduces the computations needed to maintain the LSDB. Routing within the area is determined only by the topology of the area.

The three types of routers defined by OSPF are as follows:

- Internal Router (IR): An internal router has all of its interfaces within the same area.

- Area Border Router (ABR): An ABR has interfaces in multiple areas. It is responsible for exchanging summary advertisements with other ABRs. You can create a maximum of 7 non-zero areas.
- Autonomous System Border Router (ASBR): An ASBR acts as a gateway between OSPF and other routing protocols, or other autonomous systems.

Area 0

Any OSPF network that contains more than one area is required to have an area configured as area 0, also called the backbone. All areas in an autonomous system must be connected to the backbone. When designing networks, you should start with area 0, and then expand into other areas.

The backbone allows summary information to be exchanged between ABRs. Every ABR hears the area summaries from all other ABRs. The ABR then forms a picture of the distance to all networks outside of its area by examining the collected advertisements, and adding in the backbone distance to each advertising router.

If this is the first instance of the OSPF area being used, create the area using this command:

```
create ospf area <areaid>
```

When a VLAN is configured to run OSPF, configure the area for the VLAN. If you want to configure the VLAN to be part of a different OSPF area, use this command:

```
configure ospf vlan <name> area <areaid>
```

Stub Areas

OSPF allows certain areas to be configured as stub areas. A stub area is connected to only one other area. The area that connects to a stub area can be the backbone area. External route information is not distributed into stub areas. Stub areas are used to reduce memory consumption and computation requirements on OSPF routers.

Not-So-Stubby-Areas (NSSAs)

NSSAs are similar to the existing OSPF stub area configuration option, but have two additional capabilities:

- External routes originating from an ASBR connected to the NSSA can be advertised within the NSSA.
- External routes originating from the NSSA can be propagated to other areas, including the backbone area.

The command to control the NSSA function is similar to the command used for configuring a stub area, as follows:

```
configure ospf area <area_id> nssa {summary |
nosummary} stub-default-cost <cost> {translate}
```

The `translate` option determines whether type 7 LSAs are translated into type 5 LSAs. When configuring an OSPF area as an NSSA, `translate` should only be used on NSSA border routers, where translation is to be enforced. If `translate` is not used on any NSSA border router in a NSSA, one of the ABRs for that NSSA is elected to perform translation (as indicated in the NSSA specification). Using this option on NSSA internal routers inhibits correct operation of the election algorithm.

Normal Area

A normal area is an area that is not any of:

- Area 0
- Stub area
- NSSA

Virtual links can be configured through normal areas. External routes can be distributed into normal areas.

Virtual Links

When a new area is introduced that does not have a direct physical attachment to the backbone, a virtual link is used. A virtual link provides a logical path between the ABR of the disconnected area and the ABR of the normal area that connects to the backbone. A virtual link must be established between two ABRs that have a common area, with one ABR connected to the backbone. Figure 13.1 illustrates a virtual link.

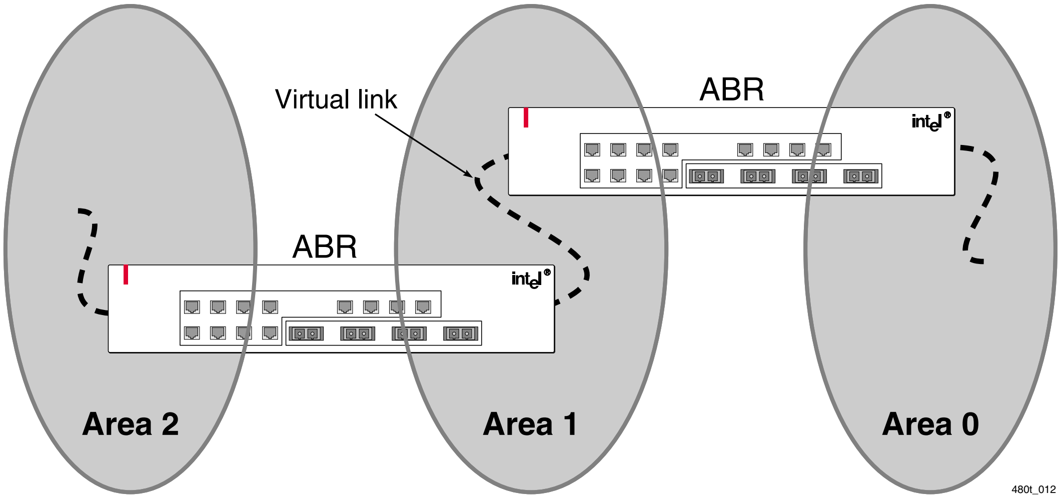

Figure 13.1: Virtual link for stub area

You can use virtual links to repair a discontiguous backbone area. In Figure 13.2, if the connection between ABR1 and the backbone fails, the ABR2 connection provides redundancy so the discontiguous area continues to communicate with the backbone using the virtual link.

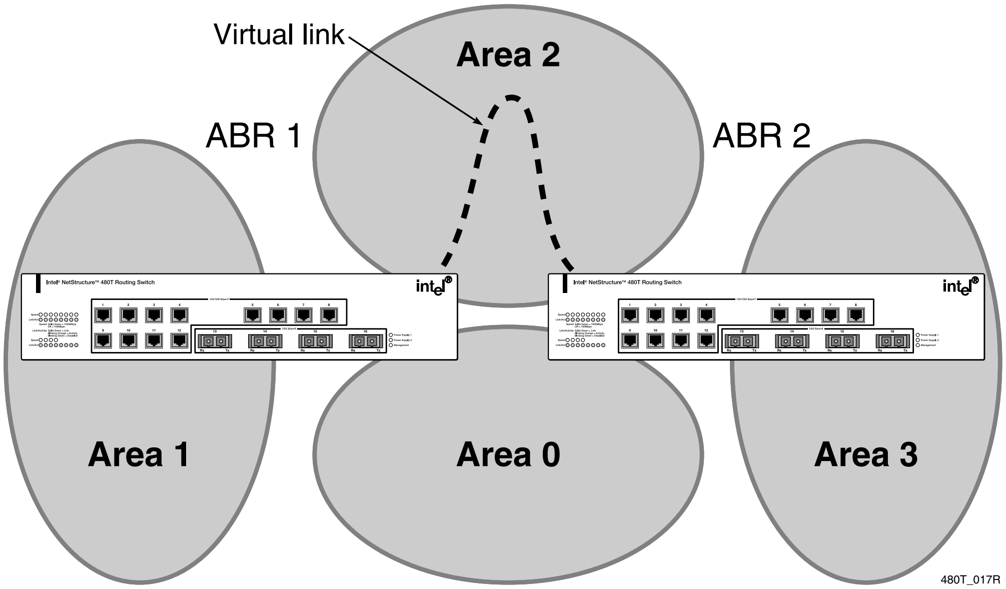

Figure 13.2: Virtual link providing redundancy

OSPF Database Overflow

The OSPF Database Overflow feature allows you to both limit the size of the LSDB and maintain a consistent LSDB across all the routers in the system.

Maintaining a consistent LSDB across all the routers in the domain ensures that all routers have a consistent view of the network.

Maintain consistency by:

- Limiting the number of external LSAs in the database of each router
- Ensuring that all routers have identical LSAs

Use this command to configure OSPF Database Overflow:

```
configure ospf ase-limit <number> {timeout
<seconds>}
```

The command allows two parameters:

- A limit specifying the number of external LSAs (excluding the default LSAs) that the system will support before it goes into overflow state. A limit value of 0 disables this functionality.
- The timeout in seconds after which the system will come out of overflow state. A timeout value of 0 leaves the system in overflow state until OSPF is disabled and enabled.

When the LSDB size limit is reached, OSPF database overflow flushes LSAs from the LSDB. OSPF database overflow flushes the same LSAs from all the routers, thereby maintaining consistency.

OSPF Passive Interface

You can configure an OSPF interface as passive. Passive interfaces do not send routing updates but do participate in receiving routing updates.

To configure an OSPF interface as a passive interface:

```
configure ospf add vlan <vlan name> area <area
identifier> {passive}
```

To reconfigure an OSPF interface as a normal interface:

```
configure ospf add vlan <vlan name> area <area
identifier>
```

To display passive interface configuration:

```
show ospf interfaces [detail]
```

Routing with OSPF

Set the RouterID

We recommend manually setting the routerID of the switches participating in OSPF instead of having the switch automatically choose its routerID based on the highest interface IP address.

Not performing this configuration in larger, dynamic environments could result in an older link state database being used. The command is:

```
configure ospf routerid <address>
```

The address is provided in dotted decimal notation. Each switch must have a unique routerID.

Route Redistribution

Both RIP and OSPF can be enabled simultaneously on the 480T routing switch. Route redistribution allows the switch to exchange routes, including static routes, between the two routing protocols. Figure 13.3 shows an example of route redistribution between an OSPF autonomous system and a RIP autonomous system.

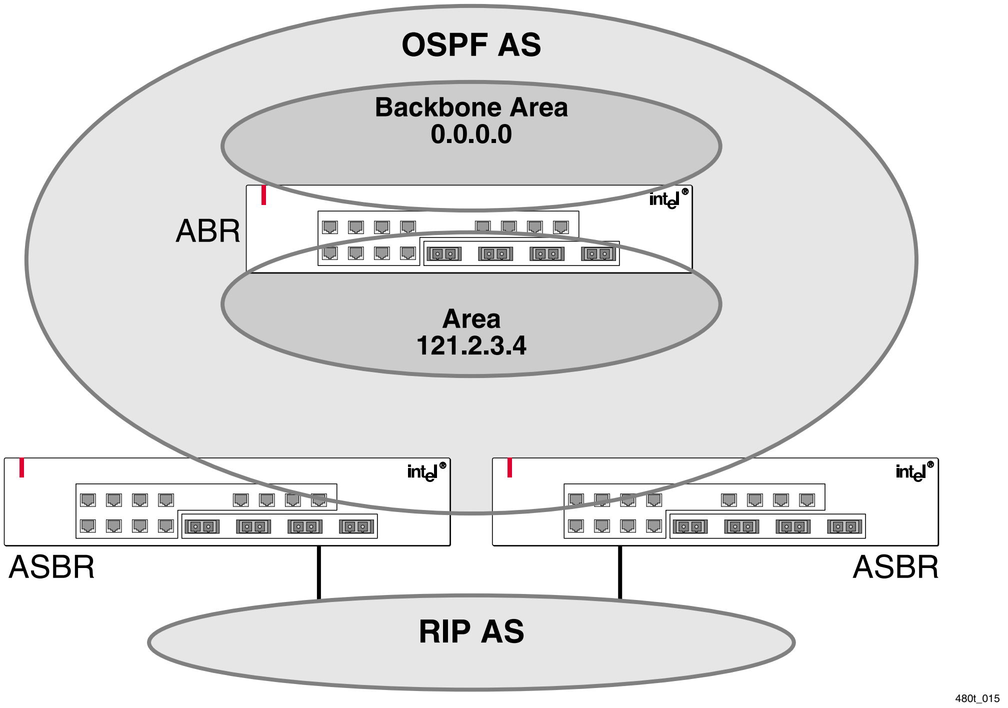

Figure 13.3: Route redistribution

Configuring Route Redistribution

Exporting routes from OSPF to RIP, and from RIP to OSPF, are discrete configuration functions. To run OSPF and RIP simultaneously, first configure both protocols, and then verify the independent operation of each. Then you can configure the routes to export from OSPF to RIP and from RIP to OSPF.

Redistributing Routes into OSPF

Use these commands to enable or disable the exporting of RIP learned, static, and direct routes to OSPF:

```
enable ospf export [static | direct | ospf | vip]
cost <metric> {tag <number>}
```

```
disable ospf export [static | direct | ospf | vip]
```

These commands enable or disable the exporting of RIP, static, and direct routes by way of LSAs to other OSPF routers as Autonomous System (AS)-external type 1 or type 2 routes. The default setting is disabled.

The cost metric is inserted for all RIP-learned, static, and direct routes injected into OSPF. If the cost metric is set to 0, the cost is inserted from the route. The tag value is used only by special routing applications. Use the number 0 if you do not have specific requirements for using a tag. The tag value in this instance has no relationship with 802.1Q VLAN tagging.

Enable or disable the export of Virtual IP addresses to other OSPF routers, using these commands:

```
enable ospf export vip cost <metric> [ase-type-1 |
ase-type-2] {tag <number>}
```

```
disable ospf export vip
```

Verify the configuration using the command:

```
show ospf
```

Redistributing Routes into RIP

Enable or disable the exporting of static, direct, and OSPF-learned routes into the RIP domain, using these commands:

```
enable rip export [static | direct | ospf | ospf-
intra | ospf-inter | ospf-extern1 | ospf-extern2 |
vip] cost <metric> tag <number>}
```

```
disable rip export [static | direct | ospf | ospf-
intra | ospf-inter | ospf-extern1 | ospf-extern2 |
vip]
```

These commands enable or disable the exporting of static, direct, and OSPF-learned routes into the RIP domain. You can choose which types of OSPF routes are injected, or you can simply choose **`ospf,`** which will inject all learned OSPF routes regardless of type. The default setting is disabled.

OSPF Timers and Authentication

Configuring OSPF timers and authentication on a per-area basis is a shorthand for applying the timers and authentication to each VLAN in the area at the time of configuration. If you add more VLANs to the area, be sure to configure the timers and authentication for the new VLANs explicitly.

OSPF Password Encryption

The neighbor password for OSPF is encrypted in upload/download configuration.

Route Map Support

See "Access Policies" on page 309.

The 480T routing switch includes the ability to apply route maps to routes that are being added to the kernel route table. You can configure the route maps based on these origins of the route:

- Direct
- Static
- RIP
- OSPF
- BGP

These route maps match the various characteristics of the route based on the originating protocol and set the accounting indexes. Use this command to configure route maps:

```
configure iproute route-map [bgp | direct | e-bgp |
i-bgp | ospf | ospf-extern1 | ospf-extern2 | ospf-
inter | ospf-intra | rip | static] [<route map> |
none]
```

Use this command to view the log:

```
show iproute {priority | vlan <vlan> | permanent |
summary | <ipaddress> <netmask> | route-map | origin
[direct | bgp | e-bgp | i-bpg | static | blackhole
| rip | ospf-intra | ospf-inter | ospf-extern1 |
ospf-extern2]} {sorted}
```

You can make dynamic changes to the route map. Direct and Static route changes are reflected immediately, while RIP, OSPF, and BGP changes are reflected within 30 seconds.

Route Map Support for OSPF Export

When OSPF is enabled the route map is applied on each and every route exported to OSPF. It can be used for filtering or for setting the cost, cost type, and tag of the exported route. You can use this feature to make dynamic changes to the route map.

Use these commands to enable OSPF route map export:

```
enable ospf export direct [[cost <metric> [ase-type-1 | ase-type-2] {tag <number>}] | <route map>]
```

```
enable ospf export static [[cost <metric> [ase-type-1 | ase-type-2] {tag <number>}] | <route map>]
```

```
enable ospf export rip [[cost <metric> [ase-type-1 | ase-type-2] {tag <number>}] | <route map>]
```

```
enable ospf export [bgp | i-bgp | e-bgp] [cost <metric> [ase-type-1 | ase-type-2] {tag <number>} | <route map>]
```

```
enable ospf export vip [[cost <metric> [ase-type-1 | ase-type-2] {tag <number>}] | <route map>]
```

BGP and OSPF Route Map Support for Tagging

Tagging support for BGP and OSPF allows you to redistribute OSPF routes from the kernel routing table to BGP, or BGP routes to OSPF.

Use this command to enable tagging:

```
configure route-map <route-map> <sequence number> [add | delete] match [nlri-list <access-profile> | as-path [access-profile <access-profile> | <as no>] | community [access-profile <access-profile> | <as no>: <number> | number <community> | no-advertise | no-export | no-export-subconfed] | next-hop <ip address> | med <number> | origin [igp | egp | incomplete] | tag <number>
```

BGP and OSPF Route Map Support for DSB Accounting

Route map support for BGP and OSPF accounting allows you to set the cost and type of the exported routes.

Use this command to enable accounting:

```
configure route-map <route-map> <sequence number>
[add | delete] set [as-path <as no> | community
[[access-profile <access-profile> | <as no>:
<number> | number <community> | no-advertise | no-
export | no-export-subconfed] | remove | [add |
delete] [access-profile <access-profile> | <as no>:
<number> | number <community> | no-advertise | no-
export | no-export-subconfed] |] | next-hop <ip
address> | med <number> | local-preference <number>
| weight <number> | origin [igp | egp | incomplete]
| tag <number> | accounting index <number> value
<number> | cost <number> | cost-type [ase-type-1 |
ase-type-2]]
```

Configuring RIP

Table 13.2 describes the commands used to configure RIP. Press the Tab key, in the command line interface, for further command options.

Table 13.2: RIP Configuration Commands

Command	Description
configure rip add vlan [<name> | all]	Configures RIP on an IP interface. For each IP interface created, RIP is disabled by default.
configure rip delete vlan [<name> | all]	Disables RIP on an IP interface. When RIP is disabled on the interface, the parameters are not reset to their defaults.
configure rip garbagetime <seconds>	Configures the RIP garbage time. The timer granularity is **10**. The default setting is **120**.
configure rip routetimeout <seconds>	Configures the route timeout. The default setting is **180**.

Table 13.2: RIP Configuration Commands (continued)

Command	Description
configure rip Rxmode [none | v1only | v2only | any] {vlan [<name> | all]}	Changes the RIP receive mode for one or all VLANs. Specify: • **none**—Drop all received RIP packets. • **v1only**—Accept only RIP v1 format packets. • **v2only**—Accept only RIP v2 format packets. • **any**—Accept both RIP v1 and v2 packets. If no VLAN is specified, the setting is applied to all VLANs. The default setting is **any**.
configure rip txmode [none | v1only | v1comp | v2only] {vlan [<name> | all]}	Changes the RIP transmission mode for one or all VLANs. Specify: • **none**—Do not transmit any packets on this interface. • **v1only**—Transmit RIP v1 format packets to the broadcast address. • **v1comp**—Transmit RIP v2 format packets to the broadcast address. • **v2only**—Transmit RIP v2 format packets to the RIP multicast address. If no VLAN is specified, the setting is applied to all VLANs. The default setting is **v2only**.
configure rip updatetime {seconds}	Changes the periodic RIP update timer. The default setting is **30**.
configure rip vlan [<name> | all] cost <number>	Configures the cost (metric) of the interface. The default setting is **1**.
enable rip	Enables RIP. The default setting is disabled.

Table 13.2: RIP Configuration Commands (continued)

Command	Description
enable rip aggregation	Enables aggregation of subnet information on interfaces configured to send RIP v2 or RIP v2-compatible traffic. The 480T routing switch summarizes subnet routes to the nearest class network route. These rules apply when using RIP aggregation: • Subnet routes are aggregated to the nearest class network route when crossing a class boundary. • Within a class boundary, no routes are aggregated. • If aggregation is enabled, the behavior is the same as in RIP v1. • If aggregation is disabled, subnet routes are never aggregated, even when crossing a class boundary. The default setting is disabled.
enable rip export [static | direct | ospf | ospf-intra | ospf-inter | ospf-extern1 | ospf-extern2 | static | vip] cost <metric> {tag <number>}	Enables RIP to redistribute routes from other routing functions. Specify: • **`static`**—Static routes • **`direct`**—Interface routes (only interfaces with IP forwarding enabled are exported) • **`ospf`**—All OSPF routes • **`ospf-intra`**—OSPF intra-area routes • **`ospf-inter`**—OSPF inter-area routes • **`ospf-extern1`**—OSPF network-unreachable route type 1 • **`ospf-extern2`**—OSPF network-unreachable route type 2 • **`vip`**—Virtual IP The cost (**`metric`**) range is 0-15. If set to `0`, RIP uses the route metric from the route origin.

Table 13.2: RIP Configuration Commands (continued)

Command	Description
enable rip originate-default {always} cost <metric> {tag <number>}	Configures a default route to be advertised by RIP if no other default route is advertised. If **`always`** is specified, RIP always advertises the default route to its neighbors. If **`always`** is not specified, RIP adds a default route if there is a reachable default route in the route table.
enable rip poisonreverse	Enables the poison reverse algorithm for RIP. The default setting is enabled. If you enable poison reverse and split horizon, poison reverse takes precedence.
enable rip splithorizon	Enables the split horizon algorithm for RIP. Default setting is enabled.
enable rip triggerupdates	Enables triggered updates. Triggered updates are a mechanism for immediately notifying a router's neighbors when the router adds or deletes routes or changes the metric of a route. The default setting is enabled.

RIP Configuration Example

Figure 13.4 illustrates a switch that has three VLANs defined as follows:

Finance

- Protocol-sensitive VLAN using the IP protocol
- Ports 1 and 3 have been assigned
- IP address 192.207.35.1

Personnel

- Protocol-sensitive VLAN using the IP protocol
- Ports 2 and 4 have been assigned
- IP address 192.207.36.1

MyCompany

- Port-based VLAN
- All ports have been assigned

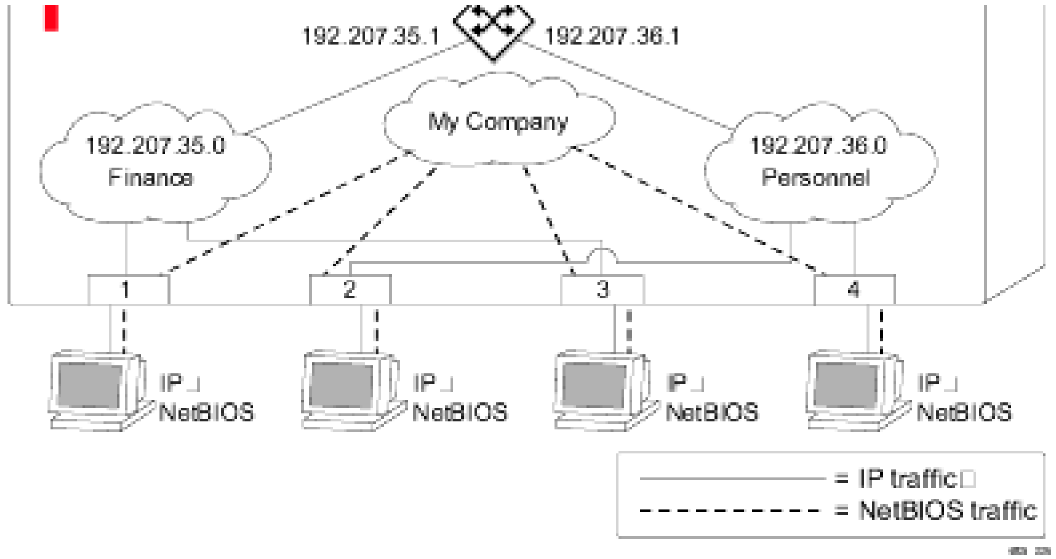

Figure 13.4: RIP configuration example

The stations connected to the system generate a combination of IP traffic and NetBIOS[§] traffic. The IP traffic is filtered by the protocol-sensitive VLANs. All other traffic is directed to the VLAN ***MyCompany***.

In this configuration, all IP traffic from stations connected to ports 1 and 3 have access to the router by way of the VLAN Finance. Ports 2 and 4 reach the router by way of the VLAN Personnel. All other traffic (NetBIOS) is part of the VLAN MyCompany.

The example in Figure 13.4 is configured as follows:

```
create vlan Finance

create vlan Personnel

create vlan MyCompany

configure Finance protocol ip
configure Personnel protocol ip
configure Finance add port 1,3
configure Personnel add port 2,4
configure MyCompany add port all

configure Finance ipaddress 192.207.35.1
configure Personnel ipaddress 192.207.36.1

enable ipforwarding
configure rip add vlan all
enable rip
```

Displaying RIP Settings

To display settings for RIP, use the commands listed in Table 13.3. For more command options, press the Tab key in the command line interface.

Table 13.3: RIP Show Commands

Command	Description
show rip {detail}	Displays RIP configuration and statistics for all VLANs.
show rip stat {detail}	Displays RIP-specific statistics for all VLANs.
show rip stat vlan <name>	Displays RIP-specific statistics for a VLAN.
show rip vlan <name>	Displays RIP configuration and statistics for a VLAN.

Resetting and Disabling RIP

To return RIP settings to their defaults, or to disable RIP, use the commands listed in Table 13.4. For more command options, press the Tab key in the command line interface.

Table 13.4: RIP Reset and Disable Commands

Command	Description
configure rip delete vlan [<name> | all]	Disables RIP on an IP interface. When RIP is disabled on the interface, the parameters are not reset to their defaults.
disable rip	Disables RIP.
disable rip aggregation	Disables the RIP aggregation of subnet information on a RIP v2 interface.
disable rip export [static | direct | ospf | ospf-intra | ospf-inter | ospf-extern1 | ospf-extern2 | vip]	Disables the distribution of non-RIP routes into the RIP domain.
disable rip originate-default	Disables the advertisement of a default route.
disable rip poisonreverse	Disables poison reverse.
disable rip splithorizon	Disables split horizon.
disable rip triggerupdates	Disables triggered updates.
unconfigure rip {vlan <name>}	Resets all RIP parameters to match the default VLAN. Does not change the enable/disable state of the RIP settings. If no VLAN is specified, all VLANs are reset.

Configuring OSPF

Each switch configured to run OSPF must have a unique routerID. We recommended manually setting the routerID of the switches participating in OSPF, instead of having each switch automatically choose its routerID based on the highest interface IP address. Not performing this configuration in larger, dynamic environments could result in an older LSDB remaining in use.

Table 13.5 describes the commands used to configure OSPF. For more command options, press the Tab key in the command line interface.

Table 13.5: OSPF Configuration Commands

Command	Description
configure ospf [area <areaid> | vlan [<name> | all]] cost [automatic | <number>]	Configures the cost metric of one or all VLAN(s). If an area is specified, the cost metric is applied to all VLANs currently within that area. When **automatic** is specified, the advertised cost is determined from the OSPF metric table and corresponds to the active highest bandwidth port in the VLAN.
configure ospf [area <areaid> | vlan [<name> | all]] priority <number>	Configures the priority used in the designated router-election algorithm for one or all IP interface(s) (VLANs) for all VLANs currently within the area. The range is 0 through 255, and the default setting is 1. Setting the value to 0 ensures that the router is never selected as the designated router or backup designated router.
configure ospf [vlan <name> | area <areaid> | virtual-link <routerid> <areaid>] authentication {encrypted} [simple-password <password> | md5 <md5_key_id> <md5_key>| none]	Specifies the authentication password (up to eight characters) or Message Digest 5 (MD5) key for one or all interfaces (VLANs) in an area. The **md5_key** is a numeric value with the range 0 to 65,536. When the OSPF area is specified, authentication information is applied to all OSPF interfaces within the area.
configure ospf add virtual-link <routerid> <areaid>	Adds a virtual link to another ABR. Specify: • **routerid**—Far-end router interface number. • **areaid**—Transit area used for connecting the two end-points.
configure ospf add vlan [<name> | all] area <areaid> {passive}	Enables OSPF on one or all VLANs (router interfaces). The **<areaid>** specifies the area to which the VLAN is assigned.
configure ospf add vlan [<name> | all] area <areaid> link-type [auto | broadcast | point-to-point]	Enables OSPF on one or all VLANs and specifies the link type.

Table 13.5: OSPF Configuration Commands (continued)

Command	**Description**
configure ospf [vlan <name> | area <areaid> | virtual-link <routerid> <areaid>] timer <retransmit_interval> <transmit delay> <hello interval> <dead interval>	Configures the timers for one interface or all interfaces in the same OSPF area. These are the default, minimum, and maximum values (in seconds): • **`retransmission_interval`** Default: `5` Minimum: `0` Maximum: `3,600` • **`transmission_delay`** Default: `1` Minimum: `0` Maximum: `3,600` • **`hello _interval`** Default: `10` Minimum: `1` Maximum: `65,535` • **`dead_interval`** Default: `40` Minimum: `1` Maximum: `2,147,483,647`
configure ospf area <areaid> add range <ipaddress> <mask> [advertise | noadvertise] [type 3 | type 7]	Configures a range of IP addresses in an OSPF area. If advertised, the range is exported as a single LSA by the ABR.
configure ospf area <areaid> delete range <ipaddress> <mask>	Deletes a range of IP addresses in an OSPF area.
configure ospf area <areaid> normal	Configures an OPSF area as a normal area. The default setting is normal.
configure ospf area <areaid> nssa [summary | nosummary] stub-default-cost <cost> {translate}	Configures an OSPF area as a NSSA.
configure ospf area <areaid> stub [summary | nosummary] stub-default-cost <cost>	Configures an OSPF area as a stub area.

Table 13.5: OSPF Configuration Commands (continued)

Command	Description
configure ospf area <areaid> [interarea-filter | external-filter] [<access-profile> | none]	Configures an OSPF area specifying filter and access profile.
configure ospf asbr-filter [<access_profile> | none]	Configures a route filter for non-OSPF routes exported into OSPF. If `none` is specified, no RIP and static routes are filtered.
configure ospf ase-limit <number> {timeout <seconds>}	Configures the OSPF Database Overflow and limits the size of the LSDB.
configure ospf ase-summary add <ipaddress> <mask> cost <cost> {tag <tag_number>}	Configures an aggregated OSPF external route using the IP addresses specified.
configure ospf ase-summary delete <ipaddress> <mask>	Deletes an aggregated OSPF external route.
configure ospf delete virtual-link <routerid> <areaid>	Removes a virtual link.
configure ospf delete vlan [<name> | all]	Disables OSPF on one or all VLANs (router interfaces).
configure ospf direct-filter [<access_profile> | none]	Configures a route filter for direct routes. If `none` is specified, all direct routes are exported, if `ospf export direct` is enabled.
configure ospf lsa-batching-interval <timer_value>	Configures the OSPF LSA batching timer value. The range is between 0 (disabled) and 600 seconds, using multiples of 5. The LSAs added to the LSDB during the interval are batched for refresh or timeout. The default setting is 30.
configure ospf metric-table 10 M <cost> 100 M <cost> 1G <cost>	Configures the automatic interface costs for 10 Mbps, 100 Mbps, and 1 Gbps interfaces; an entry is required for all three. The default cost for 10 Mbps is 10, for 100 Mbps is 5, and for 1 Gbps is 1. An entry is required for the 10 M port even though you may only need to configure the faster ports. An entry of 0 is acceptable.

Table 13.5: OSPF Configuration Commands (continued)

Command	Description
configure ospf routerid [automatic | <routerid>]	Configures the OSPF routerID. If `automatic` is specified, the 480T routing switch uses the largest IP interface address as the OSPF routerID. Manual routerID setting is recommended.
configure ospf spf-hold-time <seconds>	Configures the minimum number of seconds between Shortest Path First (SPF) recalculations. The default setting is 3.
configure ospf vlan <name> area <areaid>	Changes the area ID of an OSPF interface (VLAN).
create ospf area <areaid>	Creates an OSPF area. Area 0.0.0.0 does not need to be created. It exists by default.
disable ospf export [bgp | i-bgp | e-bgp]	Disables OSPF exporting of BGP-related routes.
enable ospf	Enables the OSPF process for the router.
enable ospf export [bgp | i-bgp | e-bgp] cost <metric> [ase-type-1 | ase-type-2] {tag <number>}	Enables OSPF to export BGP-related routes using LSAs to other OSPF routers. The default tag number is 0. The default setting is disabled.
enable ospf export direct cost <metric> [ase-type-1 | ase-type-2] {tag <number>}	Enables the distribution of local interface (direct) routes into the OSPF domain. After it is enabled, the OSPF router is considered to be an ASBR. The default tag number is 0. The default setting is disabled. If an interface route corresponds to the interface that has OSPF enabled, it is ignored.
enable ospf export rip cost <metric> [ase-type-1 | ase-type-2] {tag <number>}	Enables the distribution of RIP routes into the OSPF domain. Once enabled, the OSPF router is considered to be an ASBR. The default tag number is 0. The default setting is disabled.

Table 13.5: OSPF Configuration Commands (continued)

Command	Description
enable ospf export static cost <metric> [ase-type-1 | ase-type-2] {tag <number>}	Enables the distribution of static routes into the OSPF domain. Once enabled, the OSPF router is considered to be an ASBR. The default tag number is 0. The default setting is disabled.
enable ospf export vip cost <metric> [ase-type-1 | ase-type-2] {tag <number>}	Enables the distribution of virtual IP addresses into the OSPF domain. The default tag number is 0. The default setting is disabled.

OSPF Configuration Example

Figure 13.5 shows an example of an autonomous system using OSPF routers. The details of this network follow.

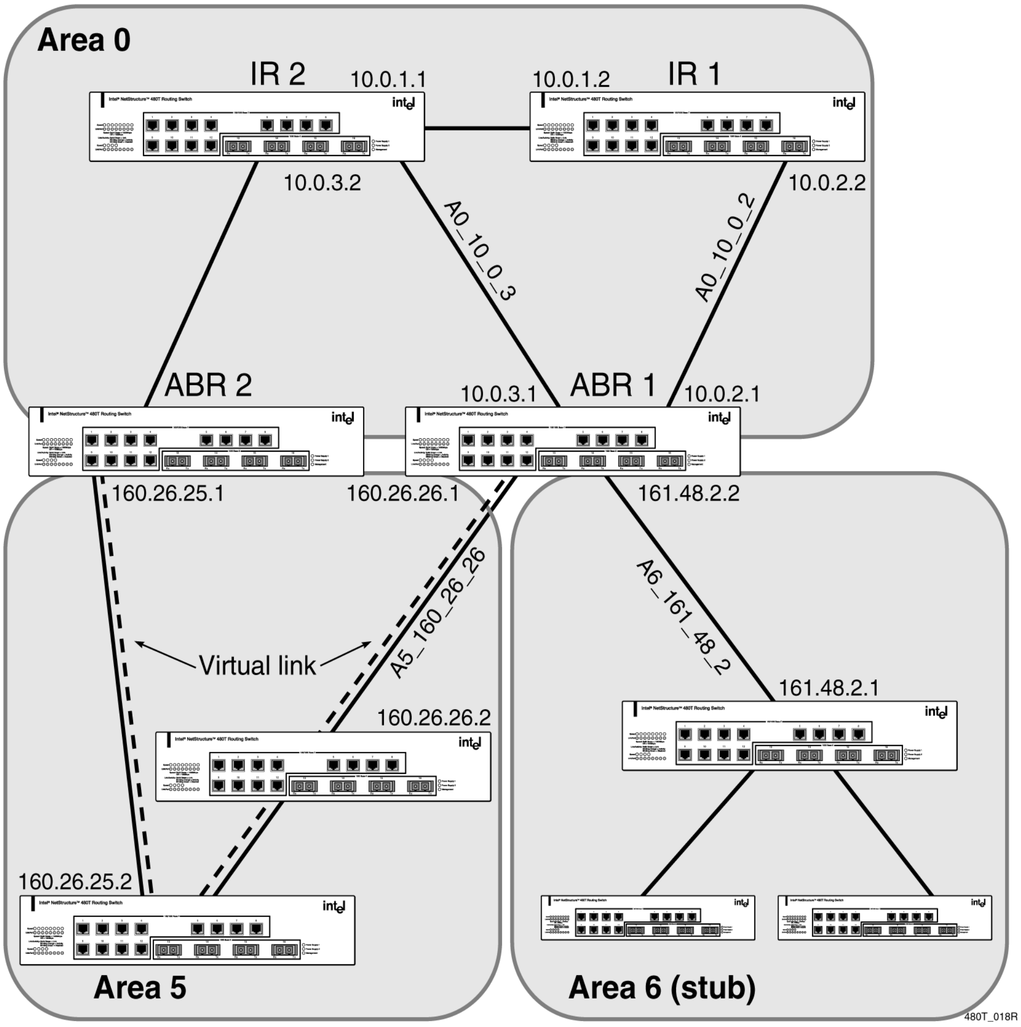

Figure 13.5: OSPF configuration example

Area 0 is the backbone area and has these characteristics:

- 2 internal routers (IR1 and IR2)
- 2 area border routers (ABR1 and ABR2)

- Network number 10.0.x.x
- 2 identified VLANs (A0_10_0_2 and A0_10_0_3)

Area 5 is connected to the backbone area through ABR1 and ABR2, having these characteristics:

- Network number 160.26.x.x
- 1 identified VLAN (A5_160_26_26)
- 2 internal routers
- A virtual link from ABR1 to ABR2 that traverses both internal routers.

  In the event that the link between either ABR and the backbone fails, the virtual link provides a connection for all routers that become discontiguous from the backbone.

Area 6 is a stub area connected to the backbone through ABR1, having these characteristics:

- Network number 161.48.x.x
- 1 identified VLAN (A6_161_48_2)
- 3 internal routers
- Uses default routes for inter-area routing

Here are two router configuration examples for Figure 13.5.

Configuration for ABR1

The following is the configuration for the router labeled ABR1:

```
create vlan A0_10_0_2

create vlan A0_10_0_3

create vlan A6_161_48_2

create vlan A5_160_26_26

configure vlan A0_10_0_2 ipaddress 10.0.2.1
255.255.255.0

configure vlan A0_10_0_3 ipaddress 10.0.3.1
255.255.255.0

configure vlan A6_161_48_2 ipaddress 161.48.2.2
255.255.255.0
```

```
configure vlan A5_160_26_26 ipaddress 160.26.26.1
255.255.255.0

create ospf area 0.0.0.5

create ospf area 0.0.0.6

enable ipforwarding

configure ospf area 0.0.0.6 stub nosummary stub-
default-cost 10

configure ospf vlan A6_161_48_2 area 0.0.0.6

configure ospf vlan A5_160_26_26 area 0.0.0.5

configure ospf add virtual-link 160.26.25.1 0.0.0.5

configure ospf add vlan all

enable ospf
```

Configuration for IR1

The following is the configuration for the router labeled IR1:

```
configure vlan A0_10_0_1 ipaddress 10.0.1.2
255.255.255.0

configure vlan A0_10_0_2 ipaddress 10.0.2.2
255.255.255.0

configure ospf add vlan all

enable ipforwarding

enable ospf
```

Displaying OSPF Settings

To display settings for OSPF, use the commands listed in Table 13.6. For more command options, press the Tab key in the command line interface.

Table 13.6: OSPF Show Commands

Command	Description
show ospf	Displays global OSPF information.
show ospf area {detail}	Displays information about all OSPF areas.
show ospf area <areaid>	Displays information about a particular OSPF area.
show ospf ase-summary	Displays the OSPF external route aggregation configuration.
show ospf interfaces {detail}	Displays information about all OSPF interfaces.
show ospf interfaces {vlan <name> | area <areaid>}	Displays information about one or all OSPF interfaces.
show ospf lsdb {detail | stats} area [<areaid> | all] [router | network | summary-net | summary-asb | network-unreachable | external-type7 | all | lsid | lstype | <routerid>]	Displays a table of the current LSDB. You can filter the display using the area ID and LSA type. The default setting is `all` with no detail. If `detail` is specified, each entry includes complete LSA information.
show ospf virtual-link [routerid <routerid> <areaid>]	Displays virtual link information about a particular router or all routers.

Resetting and Disabling OSPF Settings

To return OSPF settings to their defaults, use the commands listed in Table 13.7. For more command options, press the Tab key in the command line interface.

Table 13.7: OSPF Reset and Disable Commands

Command	Description
delete ospf area [<areaid> | all]	Deletes an OSPF area. Once removed, the associated OSPF area and interface information are deleted. Neither a backbone area, nor a non-empty area can be deleted.
disable ospf	Disables OSPF process in the router.
disable ospf export direct	Disables exporting of local interface (direct) routes into the OSPF domain.
disable ospf export rip	Disables exporting of RIP routes into the OSPF domain.
disable ospf export static	Disables exporting of statically configured routes into the OSPF domain.
disable ospf export vip	Disables exporting of virtual IP addresses into the OSPF domain.
disable ospf export [bgp | e-bgp | i-bgp]	Disables exporting of BGP, e-BGP, or i-BGP routes into the OSPF domain.
unconfigure ospf {vlan <name> | <areaid>}	Resets one or all OSPF interfaces to default settings.

14 Border Gateway Protocol (BGP)

This chapter describes how to configure the Border Gateway Protocol (BGP), an exterior routing protocol available on the Intel® NetStructure™ 480T routing switch.

For IEEE standards information, refer to http://standards.ieee.org

For more information on BGP, refer to these documents:

- RFC 1771 – *Border Gateway Protocol version 4 (BGP-4)*
- RFC 1965 – *Autonomous System Confederations for BGP*
- RFC 1966 – *BGP Route Reflection*
- RFC 1997 – *BGP Communities Attribute*
- RFC 1745 – *BGP/OSPF Interaction*

Overview

The 480T routing switch supports BGP version 4 only.

BGP is an exterior routing protocol for use in TCP/IP networks. The primary function of BGP is to allow different autonomous systems (ASs) to exchange network reachability information.

An autonomous system is a set of routers that are under a single technical administration. This set of routers uses a different routing protocol, such as Open Shortest Path First (OSPF), for intra-AS routing. One or more routers in the AS are configured as border routers, exchanging information with other border routers (in different autonomous systems) on behalf of all of the intra-AS routers.

You can use BGP as an exterior border gateway protocol (EBGP), or you can use it within an AS, as an interior border gateway protocol (IBGP).

BGP Attributes

These well-known BGP attributes are supported by the 480T routing switch:

- Origin – Defines the origin of the route. Possible values are IGP, EGP, and incomplete.
- AS_Path – The list of ASs that are traversed for this route.
- Next_hop – The IP address of the next-hop BGP router to reach the destination listed in the Network Layer Reachability Information (NLRI) field.
- Multi_Exit_Discriminator (MED) – Used to select a particular border router in another AS when multiple border routers exist.
- Local_Preference – Used to advertise this router's degree of preference to other routers within the AS.
- Atomic_aggregate – Indicates that the sending border router is using a route-aggregate prefix in the route update.
- Aggregator – Identifies the BGP router AS number and IP address that performed route aggregation.
- Community – Identifies a group of destinations that share one or more common attributes.
- Cluster_ID – Specifies a 4-byte field used by a route reflector to recognize updates from other route reflectors in the same cluster.

BGP Communities

A BGP community is a group of Border Gateway Protocol destinations that require common handling. The 480T routing switch supports these well-known BGP community attributes:

- no-export
- no-advertise
- internet

BGP Features

The BGP features supported by the 480T routing switch include:

- Route Reflectors
- Route Confederations
- Route Aggregation
- Route Map Support
- IGP Synchronization
- Loopback Interface
- OSPF-to-BGP Route Redistribution
- BGP Peer Groups

Route Reflectors

Be certain that peer routers that are not part of the cluster are fully meshed according to the rules of BGP.

One way to overcome the difficulties of creating a fully meshed AS is to use route reflectors. Route reflectors allow a single router to serve as a central routing point for the AS or sub-AS.

A cluster is formed by the route reflector and its client routers. Figure 14.1 shows a BGP cluster, including the route reflector and its clients.

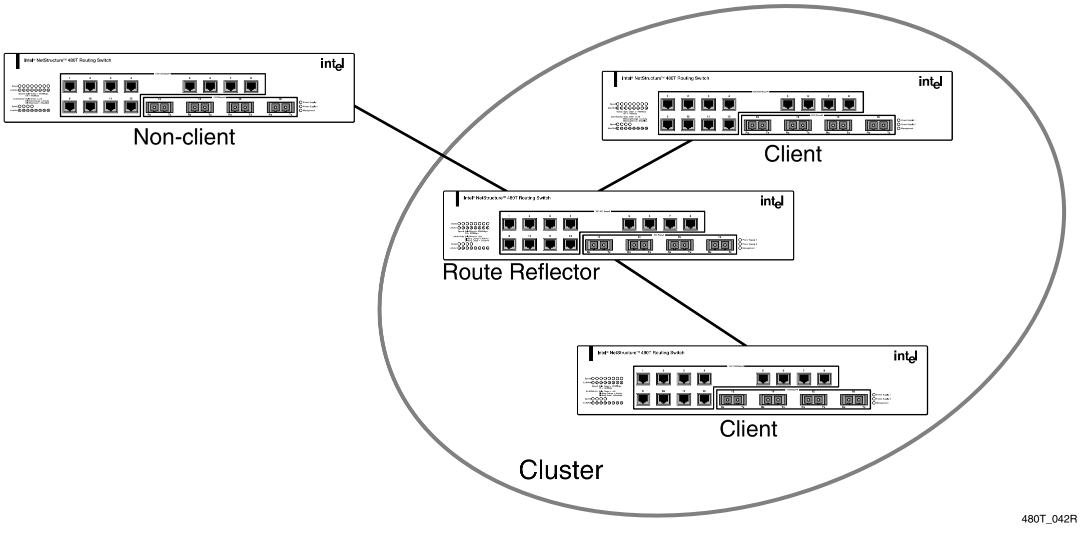

Figure 14.1: Route reflectors

Route Confederations

BGP requires networks to use a fully-meshed router configuration. This requirement does not scale well, especially when BGP is used as an interior gateway protocol.

One way to reduce the size of a fully-meshed AS is to divide the AS into multiple sub-autonomous systems and group them into a routing confederation. Within the confederation, each sub-AS must be fully meshed. The confederation is advertised to other networks as a single AS.

Route Confederation Example

Figure 14.2 shows an example of a confederation.

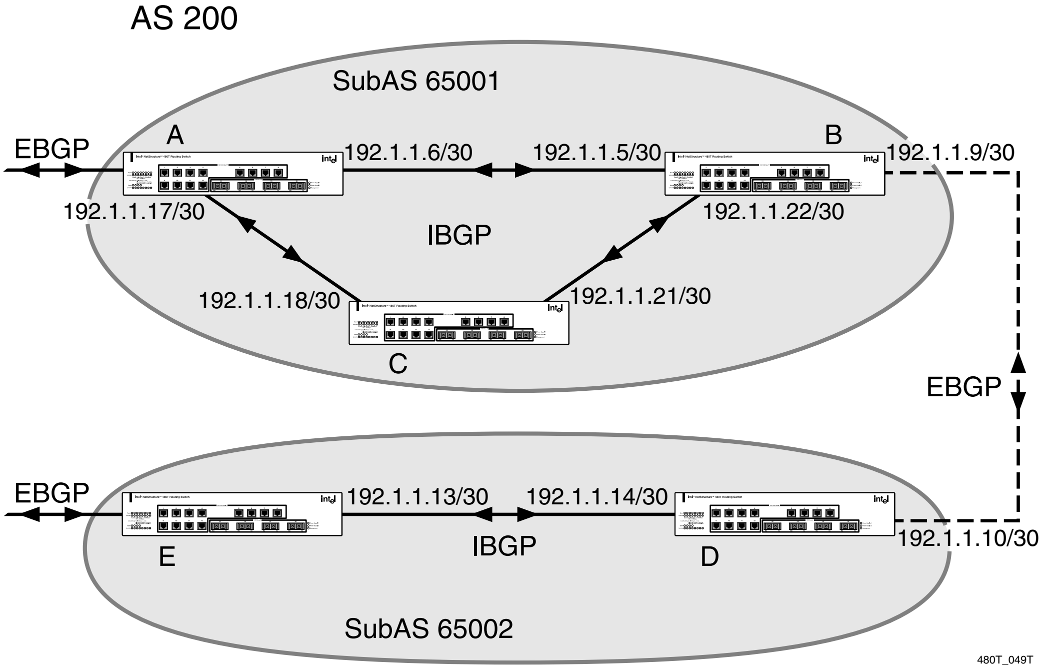

Figure 14.2: Routing confederation

In this example, AS 200 has five BGP speakers. Without a confederation, BGP would require that the routes in AS 200 be fully meshed.

Using the confederation, AS 200 is split into two sub-ASs: SubAS65001 and SubAS65002. Each sub-AS is fully meshed, and IBGP (Internal BGP) is running among its members.

EBGP (External BGP) is used between Sub65001 and SubAS65002. Router B and Router D are EBGP peers. EBGP is also used between the confederation and outside ASs.

To configure Router A, use these commands:

```
create vlan ab
configure vlan default delete port 1
configure vlan ab add port 1
configure vlan ab ipaddress 192.1.1.6/30
enable ipforwarding vlan ab
configure ospf add vlan ab area 0.0.0.0

create vlan ac
configure vlan ac add port 2
configure vlan ac ipaddress 192.1.1.17/30
enable ipforwarding vlan ac
configure ospf add vlan ac area 0.0.0.0

disable bgp
configure bgp as-number 65001
configure bgp routerid 192.1.1.17
configure bgp confederation-id 200
enable bgp

create bgp neighbor 192.1.1.5 as-number remote-AS-
number 65001
create bgp neighbor 192.1.1.18 as-number remote-AS-
number 65001
enable bgp neighbor all
```

To configure Router B, use these commands:

```
create vlan ba
configure vlan ba add port 1
configure vlan ba ipaddress 192.1.1.5/30
enable ipforwarding vlan ba
configure ospf add vlan ba area 0.0.0.0

create vlan bc
configure vlan bc add port 2
configure vlan bc ipaddress 192.1.1.22/30
enable ipforwarding vlan bc
configure ospf add vlan bc area 0.0.0.0
```

```
create vlan bd
configure vlan bd add port 3
configure vlan bd ipaddress 192.1.1.9/30
enable ipforwarding vlan bd
configure ospf add vlan bd area 0.0.0.0

disable bgp
configure bgp as-number 65001
configure bgp routerid 192.1.1.22
configure bgp confederation-id 200
enable bgp

create bgp neighbor 192.1.1.6 as-number remote-AS-
number 65001
create bgp neighbor 192.1.1.21 as-number remote-AS-
number 65001
create bgp neighbor 192.1.1.10 as-number remote-AS-
number 65002
enable bgp neighbor all
configure bgp add confederation-peer sub-AS-number
65002
```

To configure Router C, use these commands:

```
create vlan ca
configure vlan ca add port 1
configure vlan ca ipaddress 192.1.1.18/30
enable ipforwarding vlan ca
configure ospf add vlan ca area 0.0.0.0

create vlan cb
configure vlan cb add port 2
configure vlan cb ipaddress 192.1.1.21/30
enable ipforwarding vlan cb
configure ospf add vlan cb area 0.0.0.0

disable bgp
configure bgp as-number 65001
configure bgp routerid 192.1.1.21
configure bgp confederation-id 200
enable bgp

create bgp neighbor 192.1.1.22 as-number remote-AS-
number 65001
create bgp neighbor 192.1.1.17 as-number remote-AS-
number 65001
enable bgp neighbor all
```

To configure Router D, use these commands:

```
create vlan db
configure vlan db add port 1
configure vlan db ipaddress 192.1.1.10/30
enable ipforwarding vlan db
configure ospf add vlan db area 0.0.0.0

create vlan de
configure vlan de add port 2
configure vlan de ipaddress 192.1.1.14/30
enable ipforwarding vlan de
configure ospf add vlan de area 0.0.0.0

disable bgp
configure bgp as-number 65002
configure bgp routerid 192.1.1.14
configure bgp confederation-id 200
enable bgp

create bgp neighbor 192.1.1.9 as-number remote-AS-
number 65001
create bgp neighbor 192.1.1.13 as-number remote-AS
number 65002
enable bgp neighbor all
configure bgp add confederation-peer sub-AS-number
65001
```

To configure Router E, use these commands:

```
create vlan ed
configure vlan ed add port 1
configure vlan ed ipaddress 192.1.1.13/30
enable ipforwarding vlan ed
configure ospf add vlan ed area 0.0.0.0

disable bgp
configure bgp as-number 65002
configure bgp routerid 192.1.1.13
configure bgp confederation-id 200
enable bgp

create bgp neighbor 192.1.1.14 as-number remote-AS-
number 65002
enable bgp neighbor 192.1.1.14
```

Route Aggregation

Route aggregation involves combining the sub-networks of several routes so that they are advertised as a single route.

Aggregation reduces the amount of information that a BGP speaker must store and exchange with other BGP speakers. Reducing the information that is stored and exchanged also reduces the size of the routing table.

Using Route Aggregation

To use BGP route aggregation:

- Enable aggregation using this command:

  ```
  enable bgp aggregation
  ```

- Create an aggregate route, using these commands:

  ```
  configure bgp add aggregate-address <ipaddress>/
  <masklength> {as-set} {summary-only} {advertise-
  route-map <route-map>} {attribute-route-map
  <route-map>}
  ```

Route Map Support

For information see "Route Maps in BGP" on page 343.

Interior Gateway Protocol (IGP) Synchronization

You can configure an AS as a transit AS, so that it can pass traffic through from one AS to a third AS. When you configure a transit AS, it is important that the routes advertised by BGP are consistent with the routes that are available within the AS using its interior gateway protocol.

To ensure consistency, BGP should be synchronized with the IGP used within the AS. This will ensure that the routes advertised by BGP are reachable within the AS. IGP synchronization is enabled by default.

Using the Loopback Interface

If you are using BGP as your interior gateway protocol, you may decide to advertise the interface as available, regardless of the status of any particular interface. The loopback interface can also be used for EBGP multihop. Using the loopback interface eliminates multiple, unnecessary route changes.

OSPF-to-BGP Route Redistribution

You can enable both BGP and OSPF simultaneously on the 480T routing switch. Using route redistribution, the switch can exchange routes, including static routes, between the two routing protocols.

Exporting routes from OSPF to BGP, and from BGP to OSPF, are disparate configuration functions.

To run OSPF and BGP simultaneously, you must first configure both protocols and then verify the independent operation of each. Then you can configure the routes to export from OSPF to BGP and the routes to export from BGP to OSPF.

BGP Peer Groups

You can use BGP peer groups to group together up to 128 BGP neighbors. This simplifies configuring and updating neighbors because all neighbors automatically inherit the parameters of the BGP peer group.

All neighbors in the peer group share these mandatory parameters:

- Remote AS
- Source-interface
- Out-nlri-filter
- Out-aspath-filter
- Out-route-map
- Send-community
- Next-hop-self

You assign a unique name to the peer group when you create it. Use this command to create or delete a peer group.

```
[create | delete] bgp peer-group <peer-group>
```

Use these commands to configure the parameters of the peer group.

```
configure bgp peer-group <peer-group> remote-as
<number>

configure bgp peer-group <peer-group> [route-
reflector-client | no-route-reflector-client]

configure bgp peer-group <peer-group> weight
<number>

configure bgp peer-group <peer-group> source-
interface [any | vlan <vlan>]

configure bgp peer-group <peer-group> timer keep-
alive <number> hold-time <number>

configure bgp peer-group <peer-group> nlri-filter
[in | out] [none | <access profile>]

configure bgp peer-group <peer-group> as-path-
filter [in | out] [none | <access profile>]

configure bgp peer-group <peer-group> route-map-
filter [in | out] [none | <route map>]

configure bgp peer-group <peer-group> [send-
communities | dont-send-communities]

configure bgp peer-group <peer-group> soft-reset
{input | output}

configure bgp peer-group <peer-group> password
<password>

configure bgp peer-group <peer-group> [next-hop-
self | no-next-hop-self]

[enable | disable] bgp peer-group <peer-group>
soft-in-reset

[enable | disable] bgp peer-group <peer-group>
```

When you modify the parameters, the changes is applied to all neighbors in the peer group. Modifying the following parameters automatically disables and enables the neighbors before the changes take effect:

- Remote-as
- Timer
- Source-interface
- Soft-in-reset

- Password

To create a new neighbor and include it as a member of the peer group, use this command:

```
create bgp neighbor <ip address> peer-group <peer-
group> {multi-hop}
```

This creates the new neighbor as part of the peer group, and the neighbor inherits all existing parameters from the peer group. This command requires the peer group to have remote AS configured.

To add an existing neighbor to a peer group, use this command:

```
configure bgp neighbor [<ip address>| all] peer-
group <peer-group> {acquire-all}
```

If you do not specify **acquire-all**, only the mandatory parameters are inherited from the peer group. If you specify **acquire-all**, then all the parameters of the peer group are inherited. This command disables the neighbor before adding the neighbor to the peer group.

You can display existing peer groups using the command:

```
show bgp peer-group {detail | <peer-group>
{detail}}
```

If you specify **detail**, the parameters of the neighbors in the peer group that are different from the peer group are displayed.

BGP MD5 Authentication

You can configure MD5 authentication between BGP neighbors. The maximum length of the password string is 31 characters.

To configure BGP MD5 authentication:

```
configure bgp neighbor <ip address> password
<password>
```

To remove BGP MD5 authentication:

```
configure bgp neighbor <ip address> password none
```

To show BGP MD5 authentication configuration:

```
show bgp neighbor detail
```

BGP Password Encryption

The neighbor password for BGP is encrypted in upload/download configuration.

Configuring BGP

Table 14.1 describes the commands used to configure BGP. For more command options, press the Tab key in the command line interface.

Table 14.1: BGP Configuration Commands

Command	Description
configure bgp add aggregate-address <ipaddress>/<masklength> {as-set | as-match} {summary-only} {advertise-route-map <route-map>} {attribute-route-map <route-map>}	Configures an aggregate route. Options include: • **as-set**—Aggregates only the path attributes of the aggregate routes. • **summary-only**—Sends both aggregated and non-aggregated routes to the neighbors. • **advertise-route-map**—Specifies the route map used to select routes for this aggregated route. • **attribute-route-map**—Specifies the route map used to set the attributes of the aggregated route.
configure bgp add confederation-peer sub-as-number <number>	Specifies the list of sub-AS numbers that belong to a confederation. You can specify a maximum of 16 AS numbers.
configure bgp delete confederation-peer sub-AS-number <number>	Deletes a list of sub-AS numbers that belong to a confederation. You can delete up to 16 AS numbers.
configure bgp add network <ipaddress>/<mask_length> {<route_map>}	Adds a network to be originated from this router. The network must be reachable by the router.

Table 14.1: BGP Configuration Commands (continued)

Command	Description
configure bgp delete network [<ipaddress> | <mask> | all]	Deletes a network originated from this router.
configure bgp as-number <as_number>	Changes the local AS number used by BGP. *You must disable BGP before the AS number can be changed.*
configure bgp cluster-id <cluster_id>	Appends a BGP route reflector cluster-ID to the cluster list of a route. Used when multiple router reflectors are used within the same cluster of clients. *You must disable BGP before configuring the cluster ID.*
configure bgp confederation-id <confederation_id>	Changes the confederation ID.
configure bgp delete aggregate-address [<ipaddress/masklength> | all]	Deletes one or all aggregate routes.
configure bgp local-preference <local_preference>	Changes the default local-preference attribute. The range is 0 to 4,294,967,295. The default value is 100.
configure bgp med [<number> | none]	Configures the BGP multi-existence discriminator.
configure bgp neighbor [<ipaddress> | all] [route-reflector-client | no-route-reflector-client]	Configures a BGP neighbor as a route-reflector client. Implicitly defines the router as a route reflector. The neighbor must be in the same AS as the router.
configure bgp neighbor [<ipaddress> | all] [send-community | dont-send-community]	Configures whether communities should be sent to neighbors as part of the route updates. These settings apply to the peer group and all neighbors of the peer group.
configure bgp neighbor [<ipaddress> | all] [no -next-hop-self | next-hop-self]	Configures whether the next hop address used in the updates should be the address of the BGP connection originating it. These settings apply to the peer group and all neighbors of the peer group.

Table 14.1: BGP Configuration Commands (continued)

Command	Description
configure bgp neighbor [<ipaddress> | all] password [none | {encrypted} <password>]	Configures a password for a neighbor. When the password is configured, TCP MD5 authentication is enabled on the TCP connection established with the neighbor. The encrypted keyword is used in the configuration to hide the plain text password.
configure bgp neighbor [<ipaddress> | all] peer-group [none | <peer-group>]	Configures the neighbor as the member of a peer group. The acquire-all keyword is used to indicate that all parameters should be inherited from the peer group. If acquire-all is not specified, only the default parameters will be inherited from the peer group.
configure bgp neighbor [<ipaddress> | all] as-path-filter [in | out] [none | <access_profile>]	Configures an AS path filter for a neighbor. The filter is defined using the access-profile mechanism and can be installed on the input side or the output side. Use the `none` keyword to remove the filter.
configure bgp neighbor [<ipaddress> | all] nlri-filter [in | out] [none | <access_profile>]	Configures an NLRI filter for a neighbor. The filter is defined using the access-profile mechanism, and can be installed on the input side or the output side. Use the `none` keyword to remove the filter.
configure bgp neighbor [<ipaddress> | all] route-map-filter [in | out] [none | <route_map>]	Configures a route map for a neighbor. The route map can be installed on the input or output side. It is used to modify or filter the NLRI information and the path attributes associated with it, while exchanging updates with the neighbor. To remove the route map use the `none` keyword.

Table 14.1: BGP Configuration Commands (continued)

Command	Description
configure bgp neighbor [<ipaddress> | all] soft-reset {in | out}	Applies the current input or output routing policy to the routing information already exchanged with the neighbor. The input/output routing policy is determined by the **nlri-filter, as-path-filter**, and the route map configured for the neighbor in the input-output side. This command is a real-time operation and is not saved; it does not affect the 480T routing switch configuration.
configure bgp neighbor [<ipaddress> | all] source-interface [any | vlan <name>]	Changes the BGP source interface for TCP connections. The default setting is any.
configure bgp neighbor [<ipaddress> | all] timer keep-alive <seconds> hold-time <seconds>	Configures the BGP neighbor timers. The range for keep-alive is 0 to 65,535. The default keep-alive setting is 60. The range for hold-time is 0 to 21,845. The default hold-time is 90.
configure bgp neighbor [<ipaddress> | all] weight <weight>	Assigns a locally used weight to a neighbor connection for the route-selection algorithm. All routes learned from this peer are assigned the same weight. The route with the greatest weight is preferred when multiple routes are available to the same network. The range is 0 to 4,294,967,295. The default setting is 0.
configure bgp routerid <router_id>	Changes the router ID. BGP must be disabled before changing the router ID.
configure bgp soft-reconfiguration	Immediately applies the route map associated with the network command, aggregation and redistribution. This command is a real-time operation and is not saved; it does not affect the routing switch configuration.
create bgp neighbor <ipaddress> [peer-group <peergroup> | remote-as-number <as_number>] {mulithop}	Creates a new BGP peer. Use the `multihop` keyword for EBGP peers that are not directly connected.
create bgp peer-group <name>	Creates a new peer group.

Table 14.1: BGP Configuration Commands (continued)

Command	Description
disable bgp aggregation	Disables BGP route-aggregation filtering.
disable bgp always-compare-med	Disables BGP use of the Multi-Exit Discriminator (MED) from neighbors in different autonomous systems in the route-selection algorithm. MED is only used when comparing paths from the same AS. The default setting is enabled.
disable bgp export [ospf | ospf-intra | ospf-inter | ospf-extern1 | ospf-extern2] {route map}	Disables BGP from export OSPF-related routes to BGP peers.
disable bgp neighbor [<ipaddress> | all] {soft-in-reset}	Disables the soft recognition feature. Disabling the soft recognition feature can potentially limit the amount of system memory consumed by the Routing Information Base In (RIB-in).
disable bgp peer-group <peer group> {soft-in-reset}	Disables the soft recognition feature of a peer group and all the neighbors of a peer group.
disable bgp synchronization	Disables the synchronization between BGP and IGP. Default is enabled.
enable bgp	Enables BGP.
enable bgp aggregation	Enables BGP route-aggregation filtering.
enable bgp always-compare-med	Enables BGP to use the MED from neighbors in different autonomous systems in the route-selection algorithm. MED is used only when comparing paths from the same AS. The default setting is enabled.
enable bgp neighbor [<ipaddress> | all] {soft-in-reset}	Enables the BGP session. You must create the neighbor before the BGP session can be enabled.

Table 14.1: BGP Configuration Commands (continued)

Command	Description
enable bgp synchronization	Enables synchronization between BGP and IGP. When enabled, BGP waits for IGP to provide the next-hop reachability before advertising the route to an external neighbor. The default setting is enabled.
enable bgp export [ospf | ospf-intro | ospf-inter | ospf-extern1 | ospf-extern2] {<route_map>}	Configures BGP to export OSPF-related routes to BGP peers. BGP attributes associated with the OSPF routes can be applied using an optional route map.

Displaying BGP Settings

To display settings for BGP, use the commands listed in Table 14.2. For more command options, press the Tab key in the command line interface.

Table 14.2: BGP Show Commands

Command	Description
show bgp	Displays BGP configuration information.
show bgp neighbor {detail}	Disables BGP neighbor information
show bgp neighbor <ipaddress>	Displays information about a specified neighbor.

Resetting and Disabling BGP

To return BGP settings to their defaults, or to disable BGP, use the commands listed in Table 14.3. For more command options, press the Tab key in the command line interface.

Table 14.3: BGP Reset and Disable Commands

Command	Description
delete bgp neighbor [<ipaddress> | all]	Deletes one or all BGP neighbors.
disable bgp	Disables BGP.
disable bgp aggregation	Disables BGP route-aggregation.
disable bgp always-compare-med	Disables MED from being used in the route-selection algorithm.
disable bgp neighbor [<ipaddress> | all] {soft-in-reset}	Disables the BGP session. Once disabled, all the Adjacent Routing Information Base In (Adj-RIB-in) for the neighbor is flushed out.
disable bgp peer-group <peer group> {soft-in-reset}	Disables the soft recognition feature of a peer group and all the neighbors of a peer group.
disable bgp synchronization	Disables the synchronization between BGP and IGP. Default is enabled.
disable bgp export [ospf | ospf-extern | ospf-extern2 | ospf-inter | ospf-intra] <routemap>	Disables exporting OSPF routes.

BGP Route Selection

BGP will select routes based on the following precedence (from highest to lowest):

- Weight
- Local preference
- Shortest length (shortest AS path)
- Lowest origin code
- Lowest MED
- Route from external peer
- Lowest cost to next hop
- Lowest RouterID

15 IP Multicast Routing

This chapter describes the components of IP multicast routing, and how to configure it on the Intel® NetStructure™ 480T routing switch.

For more information on IP multicasting, refer to these publications:

- RFC 1112—***Host Extension for IP Multicasting***
- RFC 2236—***Internet Group Management Protocol, Version 2***
- ***DVMRP Version 3—draft_ietf_dvmrp_v3_07***
- ***PIM-DM Version 2—draft_ietf_pim_v2_dm_03***
- RFC 2326— ***Protocol Independent Multicast-Sparse Mode***

Refer to http://www.ietf.org for the Internet Engineering Task Force (IETF) Working Groups for DVMRP and PIM.

Overview

IP multicast routing allows a single IP host to send a packet to a group of IP hosts. This group of hosts can include devices that reside on the local network, within a private network, or outside of the local network.

IP multicast routing consists of these functions:

Configure IP unicast routing before you configure IP multicast routing.

- A router that can forward IP multicast packets.
- A router-to-router multicast routing protocol, for example, Distance Vector Multicast Routing Protocol (DVMRP), or Protocol Independent Multicast (PIM).
- A method for the IP host to communicate its multicast group membership to a router. For example, Internet Group Management Protocol (IGMP).

DVMRP Overview

DVMRP is a distance-vector protocol that is used to exchange routing and multicast information between routers. Like RIP, DVMRP periodically sends the entire routing table to its neighbors.

DVMRP has a mechanism that allows it to prune and graft multicast trees to reduce the bandwidth consumed by IP multicast traffic.

PIM Overview

Protocol Independent Multicast (PIM) is a multicast routing protocol similar to DVMRP. It provides both Dense Mode (PIM-DM) and Sparse Mode (PIM-SM).

The 480T routing switch supports both dense mode and sparse mode operation. You can configure dense mode or sparse mode on a per-interface basis. Once enabled, some interfaces can run dense mode, while others run sparse mode.

PIM-DM

You can run either DVMRP or PIM-DM on the switch, but not both simultaneously.

PIM-DM routers perform reverse path multicasting (RPM). However, instead of exchanging its own unicast route tables for the RPM algorithm, PIM-DM uses the existing unicast route table for the reverse path. As a result, PIM-DM requires less system memory. PIM-DM is a broadcast and prune protocol. Using PIM-DM, multicast routes are pruned and grafted in the same way as DVMRP.

PIM Sparse Mode (PIM-SM)

You can run either PIM-DM or PIM-SM on each VLAN.

Unlike PIM-DM, PIM-SM is an explicit join and prune protocol, and supports shared trees as well as shortest path trees (SPTs). The routers must explicitly be joined to one or more groups to enable communication. This is beneficial for large networks that have group members sparsely distributed.

Using PIM-SM, the router sends a join message to the rendezvous point (RP). The RP is a central multicast router that is responsible for receiving and distributing multicast packets.

When a router has a multicast packet to distribute, it encapsulates the packet in a unicast message and sends it to the RP. The RP decapsulates the multicast packet and distributes it among all member routers.

When a router determines that the multicast rate from a particular originating router (not the RP) has exceeded a configured threshold, that router can send an explicit join message to the originating router.

Once this occurs, the receiving router gets the multicast directly from the sending router, and bypasses the RP.

Static Rendezvous Points (RPs)

If you configure a static RP in your network, configure the static RP on all switches in the network.

The 480T routing switch allows you to override the PIM bootstrap message that selects a dynamic RP so that you can define a static RP in your network. To define a static RP, use the following command:

```
configure pim crp static <rp address>
```

PIM Mode Translation

A 480T routing switch functioning as a PMBR (PIM Multicast Border Router) integrates PIM-SM and PIM-DM traffic separated by the PMBR.

When forwarding PIM-DM traffic into a PIM-SM network, the PMBR will notify the RP that the PIM-DM network exists. The PMBR will then forward PIM-DM multicast packets to the RP, which will then forward the packets to those routers that have joined the multicast group.

The switch also forwards PIM-SM traffic to a PIM-DM network. The PMBR sends a join message to the RP and the PMBR then broadcasts traffic from the RP into the PIM-DM network.

There are no new commands that need to be entered to enable PIM-SM to PIM-DM functionality. By having both the DM mode interface and SM mode interface on the same router, the PMBR functionality is automatically enabled.

IP Multicast Cache Display

The `show ipmc cache` command displays a legend with a summary of each entry in the table.

IGMP Overview

IGMP is a protocol used by an IP host to register its IP multicast group membership with a router. The messaging protocol can also be snooped by a Layer 2 switch, to provide for intelligent forwarding of multicast data streams within a VLAN.

Periodically, the router queries the multicast group to see if the group is still in use. If the group is still active, a single IP host responds to the query, and group registration is maintained.

IGMP is enabled by default on the switch. However, you can configure the switch to disable the generation of periodic IGMP query packets. IGMP query should be enabled when the switch is configured to perform the IP unicast or IP multicast routing.

IGMP Snooping

IGMP snooping is a Layer 2 function that does not require multicast routing to be enabled. It reduces the flooding of IP multicast traffic.

IGMP snooping optimizes network bandwidth use, and prevents multicast traffic from flooding to parts of the network that do not need it. IP multicast traffic is not reduced in the local multicast domain.

IGMP snooping is enabled by default on the 480T routing switch. If you are using multicast routing, IGMP snooping must be enabled. If

IGMP snooping is disabled, all IGMP and IP multicast traffic will flood within a given VLAN. This is normal 802.1d bridge behavior.

IGMP and IGMP snooping must be enabled when IP unicast or multicast routing is configured (the default setting is enabled).

IGMP snooping expects at least one device in the network to periodically generate IGMP query messages. Without an IGMP querier, the switch stops forwarding IP multicast packets to the ports.

An optional optimization for IGMP snooping is the strict recognition of multicast routers only if the remote devices have joined the DVMRP (224.0.0.4) or PIM (244.0.0.13) multicast groups.

To support IGMP snooping in environments that do not have an IGMP querier, the switch can function as an IGMP querier, according to the rules of IGMP Version 2.0. If IGMP snooping is enabled, the switch periodically queries for multicast group memberships.

However, if either IGMP snooping is disabled or IGMP functionality is disabled, the switch does not generate IGMP query messages.

IGMP is enabled when the switch is configured to perform IGMP snooping and there is no other reliable querier on the network.

IGMP Leave Message

IGMP snooping supports the IGMP leave message. When a port sends an IGMP leave message, the switch removes the IGMP snooping entry after 10 seconds. The router still sends a query to determine which ports wish to remain in the multicast group. If other members of the VLAN wish to remain in the multicast group, the router will ignore the leave message, but the port is removed from the IGMP snooping table.

If the last port within a VLAN sends an IGMP leave message, the router will not receive any responses to the query, and the router will immediately remove the VLAN from its multicast group.

IGMP Display

The **`show igmp snooping`** command can be displayed with a summary or detail view.

IGMP Query Interval

The maximum value you can set for the IGMP query interval is 429,496,729. The values you can set for query response interval and the last member query interval are between 1 second and 25 seconds.

IGMP Configuration Commands

Table 15.1 describes the commands used to configure the Internet Gateway Message Protocol (IGMP). For more command options, press the Tab key in the command line interface.

Table 15.1: IGMP Configuration Commands

Command	Description
enable igmp {vlan <name>}	Enables IGMP on a router interface. If no VLAN is specified, IGMP is enabled on all router interfaces. The default setting is enabled.
enable igmp snooping {forward-mcrouter-only}	Enables IGMP snooping on the switch. If **`forward-mcrouter-only`** is specified, the switch forwards all multicast traffic to the multicast router only. Otherwise, the switch forwards all multicast traffic to any IP router.

Table 15.1: IGMP Configuration Commands (continued)

Command	Description
configure igmp <query_interval> <query_response_interval> <last_member_query_interval>	Configures the IGMP timers. Timers are based on IEEE RFC2236. Specify: • `query_interval`—The amount of time, in seconds, the system waits between sending out general queries. The range is 1 to 429,496,729 seconds. The default setting is 125. • `query_response_interval`—The maximum response time inserted into the periodic general queries. The range is 1 to 25. The default setting is 10. • `last_member_query_interval`—The maximum response time inserted into a group-specific query sent in response to a leave-group message. The range is 1 to 25. The default setting is 1.
configure igmp snooping timer <router_timeout> <host_timeout>	Configures the IGMP snooping timers. Timers should be set to approximately 2.5 times the router-query interval in use on the network. Specify: • `router_timeout`—The interval, in seconds, between the last time the router was discovered and the current time. The range is 10 to 2,147,483,647 seconds (68 years). The default setting is 260. • `host_timeout`—The interval, in seconds, between the last IGMP group report message from the host and the current time. The range is 10 to 2,147,483,647 seconds (68 years). The default setting is 260.

Configuring IP Multicasting Routing

To configure IP multicast routing:

1. Configure the system for IP unicast routing.
2. Enable multicast routing on the interface, using this command:

   ```
   enable ipmcforwarding {vlan <name>}
   ```

3. Enable DVMRP or PIM on all IP multicast routing interfaces, using either:

   ```
   configure dvmrp add vlan [<name> | all]
   configure pim add vlan [<name> | all] {dense |
   sparse}
   ```

4. Enable DVMRP or PIM on the router, using either:

   ```
   enable dvmrp
   enable pim
   ```

Table 15.2 describes the commands used to configure IP multicast routing. Press the Tab key in the command line interface for further command options.

Table 15.2: IP Multicast Routing Configuration Commands

Command	Description
enable dvmrp {[Rxmode | txmode] vlan [<name> | all]}	Enables DVMRP on the system.
configure dvmrp add vlan [<name> | all]	Enables DVMRP on one or all IP interfaces. If no VLAN is specified, DVMRP is enabled on all IP interfaces. When an IP interface is created, DVMRP is disabled by default.
configure dvmrp delete vlan [<name> | all]	Disables DVMRP on one or all IP interfaces. If no VLAN is specified, DVMRP is disabled on all IP interfaces.

Table 15.2: IP Multicast Routing Configuration Commands (continued)

Command	Description
configure dvmrp timer <route_report_interval> <route_replacement_time>	Configures the global DVMRP timers. Specify the following: • **`route_report_interval`**—how many seconds the system waits between transmitting periodic route report packets. The range is 1 to 2,147,483,647 seconds (68 years). The default setting is 60. Because triggered update is always enabled, the route report will always be transmitted prior to the expiration of the route report interval. • **`route_replacement_time`**—The hold-down time before a new route is learned, after the previous route is deleted. The range is 1 to 2,147,483,647 seconds (68 years). The default setting is 140.
configure dvmrp vlan [<name> | all] cost <number>	Configures the cost (metric) of the interface. The default setting is 1.
configure dvmrp vlan [<name> | all] export-filter [<access_profile> | <none>]	Configures DVMRP to filter out routes specified in the export filter when sending out route advertisements.
configure dvmrp vlan [<name> | all] import-filter [<access_profile> | <none>	Configures DVMRP to filter out certain routes (defined by the access profile) received from a neighbor.
configure dvmrp vlan [<name> | all] trusted-gateway [<access_profile> | <none>]f	Configures the DVMRP trusted gateway, based on the access profile.

Table 15.2: IP Multicast Routing Configuration Commands (continued)

Command	Description
configure dvmrp vlan <name> timer <probe_interval> <neighbor timeout>	Configures DVMRP interface timers. Specify: • **`probe_interval`**—How many seconds the system waits between transmitting DVMRP probe messages. The range is 1 to 2,147,483,647 seconds (68 years). The default setting is 10. • **`neighbor_timeout_interval`**—The amount of time before a DVMRP neighbor route is declared to be down. The range is 1 to 2,147,483,647 seconds (68 years). The default setting is 35.
configure pim add vlan [<vlan> | all] {dense | sparse}	Enables PIM on an IP interface. When an IP interface is created, per-interface PIM configuration is disabled by default. The default PIM mode is dense.
configure pim cbsr [vlan <name> priority <priority> | none]	Configures a candidate bootstrap router for PIM sparse-mode operation. The range is 0 - 255. The default setting is 0 and indicates the lowest priority. To delete a Candidate Bootstrap Router (CBSR), use the keyword **`none`** as the priority.

Configuration Examples

See Chapter 13, "RIP and OSPF" on page 223 for information on configuring OSPF.

In the example below, the system labeled IR1 is configured for IP multicast routing using PIM-DM. l

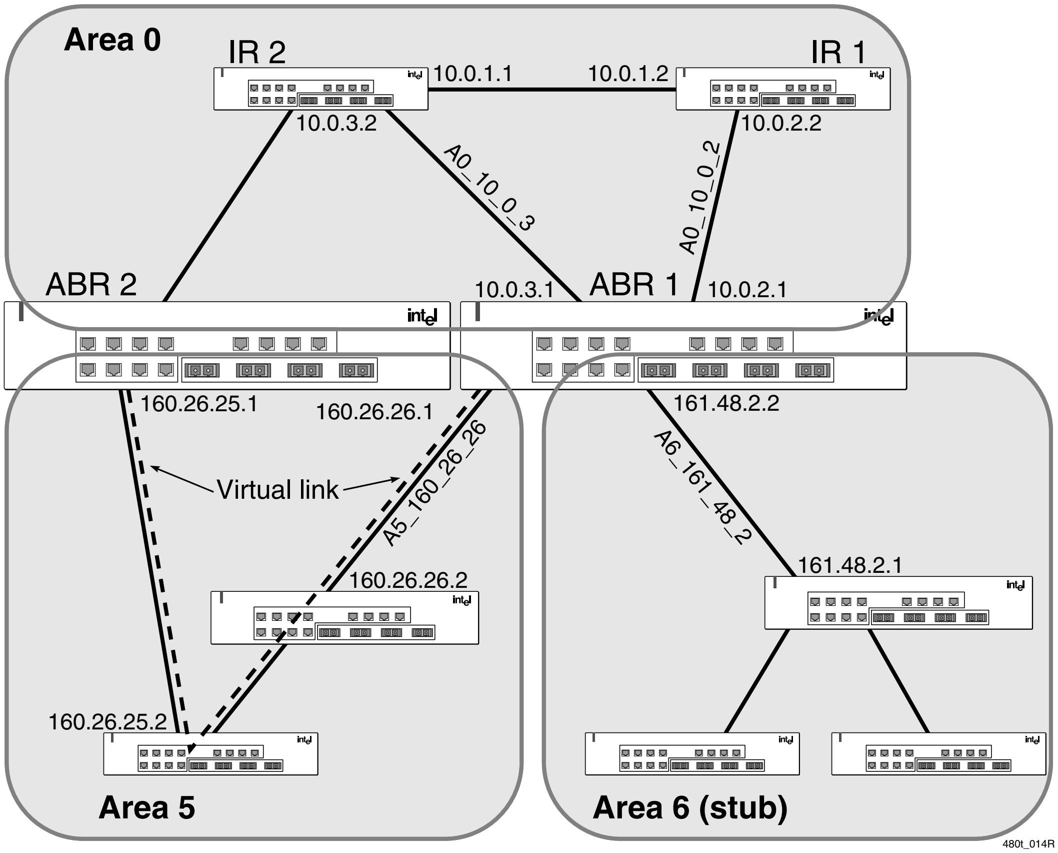

Figure 15.1: IP multicast routing PIM-DM configuration example

Configuration for IR1

The following is the configuration for the router labeled IR1:

```
configure vlan A0_10_0_1 ipaddress 10.0.1.2
255.255.255.0
```

```
configure vlan A0_10_0_2 ipaddress 10.0.2.2
255.255.255.0

configure ospf add vlan all
enable ipforwarding
enable ospf
enable ipmcforwarding
configure pim add vlan all
enable pim
```

PIM-SM Configuration Example

In this example, the system labeled ABR1 is configured for IP multicast routing using PIM-SM.

Figure 15.2: IP multicast routing using PIM-SM configuration

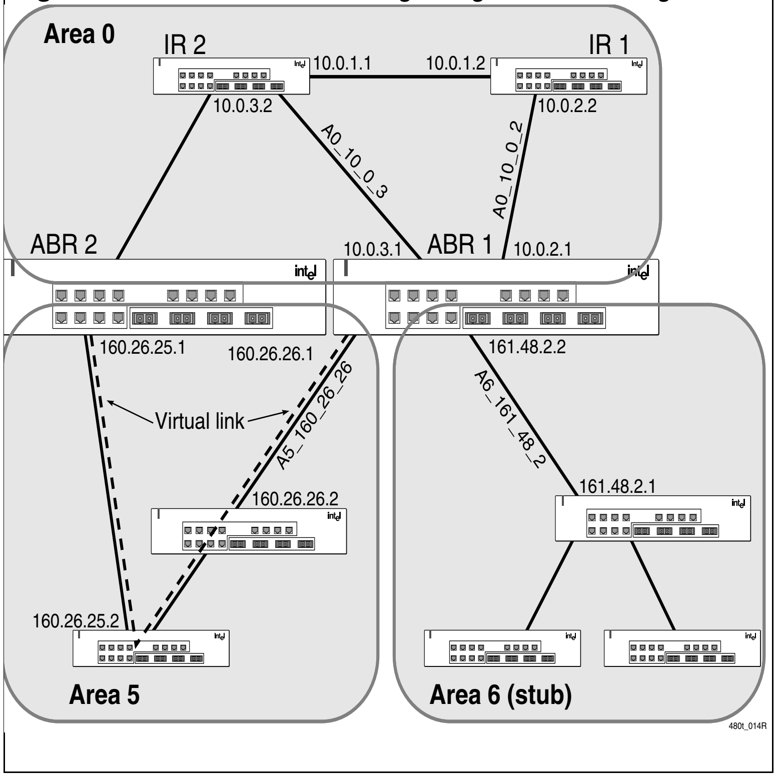

Configuration for ABR1

The following is the configuration for the router labeled ABR1:

```
configure vlan A0_10_0_2 ipaddress 10.0.2.1
255.255.255.0

configure vlan A0_10_0_3 ipaddress 10.0.3.1
255.255.255.0

configure vlan A6_161_48_2 ipaddress 161.48.2.2
255.255.255.0

configure vlan A5_160_26_26 ipaddress 160.26.26.1
255.255.255.0

configure ospf add vlan all

enable ipforwarding

enable ipmcforwarding

configure pim add vlan all sparse

create access-profile rp-list ipaddress

configure rp-list add ipaddress 224.0.0.0 240.0.0.0

enable loopback-mode A0_10_0_3

configure pim crp A0_10_0_3 rp-list 30

configure pim cbsr A0_10_0_3 30

configure pim spt-threshold 16 8
```

Displaying IP Multicast Routing Settings

To display settings for IP multicast routing components, use the commands listed in Table 15.3. For more command options, press the Tab key in the command line interface.

Table 15.3: IP Multicast Routing Show Commands

Command	Description
show dvmrp {vlan <name> | route {detail}}	Displays the DVMRP configuration and statistics, or the unicast route table. The default setting is all.

Table 15.3: IP Multicast Routing Show Commands (continued)

Command	Description
show igmp snooping {vlan <name> | detail}	Displays IGMP snooping registration information, and a summary of all IGMP timers and states.
show ipmc cache {detail} {<group>} {<src_ipaddress> <mask>}	Displays the IP multicast forwarding cache.
show pim {vlan <name> | detail}	Displays the PIM configuration and statistics. If no VLAN is specified, the configuration is displayed for all PIM interfaces.
show pim rp-set {group}	Displays the RP set for one or all groups.

Deleting and Resetting IP Multicast Settings

To return IP multicast routing settings to their defaults and disable IP multicast routing functions, use the commands listed in Table 15.4. For more command options, press the Tab key in the command line interface.

Table 15.4: IP Multicast Routing Reset and Disable Commands

Command	Description
clear igmp snooping {vlan <name>}	Removes one or all IGMP snooping entries.
clear ipmc [counters | cache | debug | trace | dlcs | fdb | igmp | iparp | ipfdb | log | session | slb] {<group> {<src_ipaddress> <mask>}}	Resets the IP multicast items. If no options are specified, all IP multicast entries are flushed.
configure ipmc cache timeout <seconds>	Configures the aging time (in seconds) for multicast cache entries. The default setting is `300`.

Table 15.4: IP Multicast Routing Reset and Disable Commands (continued)

Command	Description
disable dvmrp {[Rxmode | txmode] vlan [<name> | all]}	Disables DVMRP on the system.
disable dvmrp Rxmode vlan [<name> | all]	Disables receiving of DVMRP packets on a per-VLAN basis.
disable dvmrp txmode vlan [<name> | all]	Disables transmitting of DVMRP packets on a per-VLAN basis.
disable igmp {vlan <name>}	Disables the router-side IGMP processing on a router interface. No IGMP query is generated, but the switch continues to respond to IGMP queries received from other devices. If no VLAN is specified, IGMP is disabled on all router interfaces.
disable igmp snooping	Disables IGMP snooping. IGMP snooping can be disabled only if IP multicast routing is not being used. Disabling IGMP snooping allows all IGMP and IP multicast traffic to flood within a given VLAN.
disable ipmcforwarding {vlan <name>}	Disables IP multicast forwarding.
disable pim	Disables PIM on the system.
unconfigure dvmrp {vlan <name>}	Resets the DVMRP timers to their default settings. If no VLAN is specified, all interfaces are reset.
unconfigure igmp	Resets all IGMP settings to their default values and clears the IGMP group table.
unconfigure pim {vlan <name>}	Resets all PIM settings to their default values.

16 IPX Routing

This chapter describes how to configure IPX§, IPX/RIP, and IPX/SAP on the Intel® NetStructure™ 480T routing switch. It assumes that you are familiar with IPX. If not, refer to your Novell documentation.

Overview of IPX

The 480T routing switch provides support for IPX, IPX/RIP (Routing Information Protocol), and IPX/SAP (Service Advertisement Protocol). The switch dynamically builds and maintains an IPX routing table and an IPX service table.

The switch supports separate routing interfaces for IP and IPX traffic on the same VLAN, load sharing of IPX routed traffic, and 802.1Q tagged packets on a routed IPX VLAN.

Router Interfaces

The routing software and hardware routes IPX traffic between IPX router interfaces. A router interface is simply a VLAN that has an IPX network identifier (NetID) and IPX encapsulation type assigned to it.

As you create VLANs with different IPX NetIDs, the switch automatically routes between them. Both the VLAN switching and IPX routing function occur within the switch. You can configure a VLAN with either an

IPX NetID or an IP address. You also can configure a VLAN for both IPX and IP routing.

Figure 16.1 shows the same switch discussed earlier in Figure 12.1 on page 191. In Figure 16.1, IPX routing is added to the switch, and two additional VLANs have been defined—***Exec*** and ***Support***. Both VLANs have been configured as protocol-specific VLANs, using IPX.

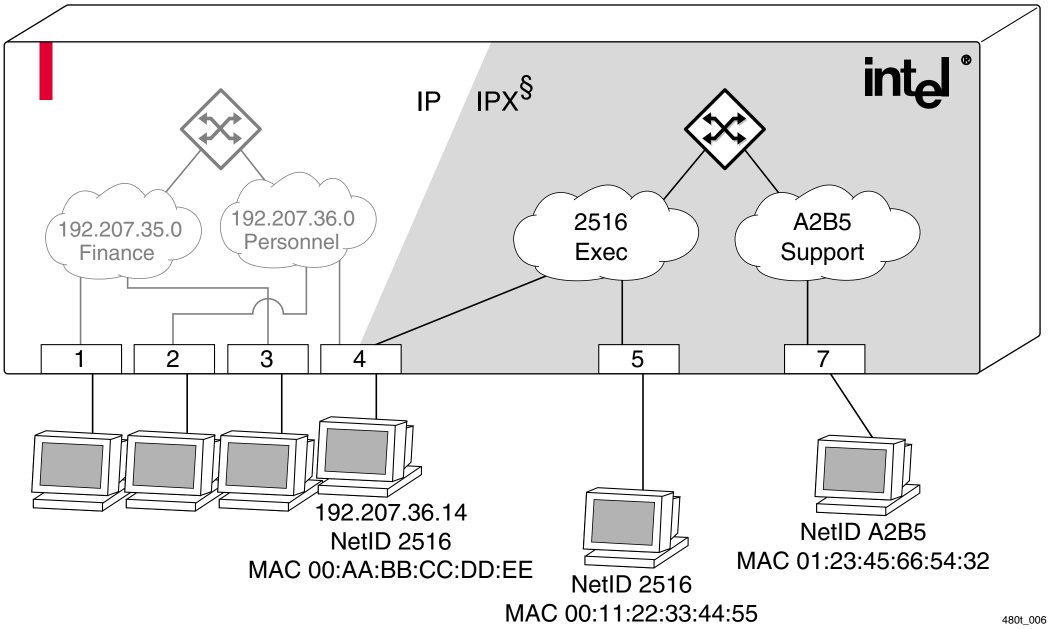

Figure 16.1: IPX VLAN configuration

Exec is assigned the IPX NetID 2516. ***Support*** is assigned the IPX NetID A2B5. Port 5 is assigned to ***Exec***; Port 7 is assigned to ***Support***. In addition, port 4 is assigned to ***Exec***. Therefore, port 4 belongs to both the ***Personnel*** VLAN (running IP) and the ***Exec*** VLAN (running IPX).

Traffic within each VLAN is switched using the Ethernet MAC address. Traffic between ***Exec*** and ***Support*** is routed using the IPX NetID. Traffic cannot be sent between the IP VLANs (***Finance*** and ***Personnel***) and the IPX VLANs (***Exec*** and ***Support***).

IPX Encapsulation Types

Novell NetWare§ supports four types of frame encapsulation. The term for each type is described in Table 16.1.

Table 16.1: IPX§ Encapsulation Types

Name	Description
ENET_II	The frame uses the Ethernet 2 header.
ENET_8023	The frame includes the IEEE 802.3 length field, but does not include the IEEE 802.2 Logical Link Control (LLC) header. This encapsulation is used by NetWare§ version 2.x and the original 3.x version.
ENET_8022	The frame uses the IEEE format and includes the IEEE 802.2 LLC header. This encapsulation is used by NetWare version 3.12 and 4.x.
ENET_SNAP	The frame adds a Subnetwork Access Protocol (SNAP) header to the IEEE 802.2 LLC header.

To configure a VLAN to use a particular encapsulation type, use this command:

```
configure vlan <name> xnetid <netid> [enet_ii |
enet_8023 | enet_8022 | enet_snap]
```

IPX and IP

The 480T routing switch supports:

- Separate routing interfaces for IP and IPX traffic on the same VLAN
- Load sharing of IPX routed traffic
- 802.1Q tagged packets on a routed IPX VLAN

IP and IPX on the Same VLAN

The switch supports IP and IPX routing within the same VLAN. This feature does not require any special configuration.

Tagged IPX VLAN

The switch supports tagged 802.1Q traffic on an IPX VLAN that is performing routing.

Tagging is most commonly used to create VLANs that span multiple switches. Using VLAN tags, multiple VLANs can span multiple switches using one or more trunks. In a port-based VLAN, each VLAN requires its own pair of trunk ports.

Another benefit of tagged VLANs is the ability to have a port be a member of multiple VLANs. This is particularly useful if you have a device (such as a server) that must belong to multiple VLANs. A single port can be a member of only one port-based VLAN. All additional VLAN memberships for that port must be configured with tags.

To configure a tagged IPX VLAN, assign a tag to the VLAN using this command:

```
configure vlan <name> tag <vlanid>
```

The valid range is from 1 to 4095.

To assign tagged ports to the VLAN, use this command:

```
configure vlan <name> add port <portlist> {tagged |
untagged} {nobroadcast}
```

To display your VLAN settings, use this command:

```
show vlan {<name>} {detail}
```

IPX Load Sharing

See "Load Sharing" on page 84.

The 480T routing switch supports IPX load sharing. There is no special configuration requirement to support this function. Simply configure load sharing as you would normally.

Populating the Routing Table

The routing switch builds and maintains an IPX routing table. As in the case of IP, the table is populated using dynamic and static entries.

Dynamic Routes

Dynamic routes are typically learned using IPX/RIP. Routers that use IPX/RIP exchange information in their routing tables in the form of advertisements. Using dynamic routes, the routing table contains only networks that are reachable.

Dynamic routes are aged out of the table when an update for the network is not received for a period of time, as determined by the routing protocol.

Static Routes

Static routes are manually entered into the routing table. Static routes are used to reach networks not advertised by routers. You can configure up to 64 static IPX routes on the 480T routing switch. Static routes are never aged out of the routing table. Static routes are advertised to the network using IPX/RIP.

IPX/RIP Routing

See Figure 13.4 on page 241.

The switch supports the use of IPX/RIP for unicast routing. IPX/RIP is different from IP/RIP. However, many of the concepts are the same. The 480T routing switch supports these IPX/RIP features:

- Split horizon
- Poison reverse
- Triggered updates

Route information is entered into the IPX route table in one of two ways:

- Dynamically, using RIP
- Statically, using the command:

```
configure ipxroute add [<dest_netid> | default]
next_hop_netid next_hop_node_addr <hops> <ticks>
```

IPX/RIP is automatically enabled when a NetID is assigned to the VLAN. To remove the advertisement of an IPX VLAN, use the command:

```
configure ipxrip delete {vlan <name> | all}
```

GNS Support

The 480T routing switch supports the Get Nearest Server (GNS) reply function. When a NetID is assigned to the switch, the GNS reply service is automatically enabled. When a station requests a particular service on the network (for example, locating a print server), the station sends a GNS request and the switch responds to the request. If GNS-reply is disabled, the switch drops the request.

To disable GNS-reply, use this command:

```
disable ipxsap gns-reply {vlan <name>}
```

Routing SAP Advertisements

The 480T routing switch contains an IPX Service Table, and propagates SAP advertisements to other IPX routers on the network. Each SAP advertisement contains:

- Service type
- Server name
- Server NetID
- Server node address

The service information is entered into the IPX Service Table in one of two ways:

- Dynamically, through SAP
- Statically, using this command:

  ```
  configure ipxservice add <service_type>
  <service_name> <netid> <node_address> <socket>
  <hops>
  ```

Configuring IPX

This section describes the commands associated with configuring IPX, IPX/RIP, and IPX/SAP on the 480T routing switch. Configure IPX routing as follows:

1. Create at least two VLANs (see "Virtual LANs (VLANs)" on page 95).
2. If you are combining an IPX VLAN with another VLAN on the same port(s), you must use a protocol filter on one of the VLANs, or use 802.1Q tagging.
3. Assign each VLAN a NetID and encapsulation type, using this command:

   ```
   configure vlan <name> xnetid <netid> [enet_ii |
   enet_8023 | enet_8022 | enet_snap]
   ```

Ensure that each VLAN has a unique IPX NetID and that the encapsulation type matches the VLAN protocol.

Once you configure the IPX VLAN information, IPX forwarding automatically begins to function. Specifically, configuring the IPX VLAN automatically enables the IPX/RIP, IPX/SAP, and SAP GNS services.

Verifying IPX Router Configuration

Use these commands to verify the IPX routing configuration:

- **`show vlan`**—Along with other information, this command displays the IPX NetID setting and encapsulation type.
- **`show ipxconfig`**—Analogous to the **`show ipconfig`** command for the IP protocol, it displays summary global IPX configuration information followed by per-VLAN information. Information includes enable/disable status for IPX/RIP, IPX/SAP, IPX route sharing, IPX service sharing, and so on.
- **`show ipxroute`**—Analogous to the **`show iproute`** command for the IP protocol. it displays static and learned routes, along with information about the VLAN that uses the route, hop count, age of the route, and so on.
- **`show ipxsap`**—Displays the enable status of IPX/SAP for the VLAN, and its operational and administrative status (including the GNS reply service). It also lists any identified IPX/SAP neighbors, SAP packet statistics, and several other timer settings.

- **show ipxrip**—Displays the enable status of IPX/RIP for the VLAN, including operational and administrative status. It also lists identified IPX/RIP neighbors, RIP packet statistics, and several other timer settings.
- **show ipxservice**—Displays the contents of the IPX Service Table.

Protocol-Based VLANs for IPX

When combining IPX VLANs with other VLANs on the same physical port, it may be necessary to assign a protocol filter to the VLAN. This is especially true if it is not possible to use 802.1Q VLAN tagging.

For convenience, IPX-specific protocol filters have been defined and named in the default configuration of the switch. Each filter is associated with a protocol encapsulation type. The IPX-specific protocol filters and the associated encapsulation type of each are described in Table 16.2.

Table 16.2: IPX§ Protocol Filters and Encapsulation Types

Protocol Name	Protocol Filter	Used for Filtering IPX Encapsulation Type
IPX	etype 0x8137	enet_ii
IPX_8022	llc 0xe0e0	enet_802_2
IPX_snap	SNAP 0x8137	enet_snap

It is not possible to define a protocol-sensitive VLAN for filtering the IPX **enet_8023** encapsulation type. Instead, use a protocol-sensitive filter on the other VLANs that share the same ports, leaving the **enet_8023** encapsulation VLAN configured using the **any** protocol.

Tuning

On larger networks, increase IPX SAP and IPX RIP update intervals to reduce CPU load (e.g., from default of 60 to 120 seconds).

To increase route stability, you can increase the hold multiplier (default is 3 for 180 seconds). To modify these parameters use CLI commands:

```
configure ipxrip <vlan name> update-interval <time>
hold-multiplier <number>
```

```
configure ipxsap <vlan name> update-interval <time>
hold-multiplier <number>
```

Tagged VLANs and IPX

IPX routing is not supported on tagged VLANs.

IPX and Round-Robin Load Sharing

Due to packet sequencing problems, we do not recommend that IPX load sharing run with the round-robin load-sharing algorithm.

IPX Performance Testing Using Traffic Generators

When using traffic generation equipment to test the wire-speed capability of IPX routing, entries that are allowed to age out with the ports remaining active cannot be re-learned on that port and will not be forwarded at wire-speed.

Restarting the port or clearing the FDB will not address this issue. In a “real-world” IPX environment, clients and servers generally do not lose communication with the directly attached switch for the FDB entries to age out.

IPX and Bi-Directional Rate Shaping

Bi-directional rate shaping is not supported for use with IPX traffic.

IPX Commands

Table 16.3 describes the commands used to configure basic IPX settings. For more command options, press the Tab key in the command line interface.

Table 16.3: Basic IPX§ Commands

Command	Description
configure ipxmaxhops <number>	Configures the IPX maximum hop count when forwarding IPX packets. The default setting is 16. Change this only if NetWare§ Link Services Protocol (NLSP) is running in the IPX network.
configure ipxroute add [<dest_netid> | default] <next_hop_id> <next_hop_node_addr> <hops> <tics>	Adds a static IPX route entry in the IPX route table. Specify: • `next_hop_id`—The NetID of the neighbor IPX network. • `next_hop_node_addr`—The node address of the next IPX router. • `hops`—The maximum hop count. • `tics`—The timer delay value. You can enter up to 64 static routes.
configure ipxroute delete [<dest_netid> | default] <next_hop_netid> <next_hop_node_addr>	Removes a static IPX route entry from the route table.
configure ipxservice add <service_type> <service_name> <netid> <node_address> <socket> <hops>	Adds a static entry to the IPX service table. Specify: • `service_type`—Srvice type. • `service_name`—Service name. • `netid`—IPX network identifier. • `node_address`—Node address of the server. • `socket`—IPX port number. • `hops`—The number of hops (for SAP routing).

Table 16.3: Basic IPX§ Commands (continued)

Command	Description
configure ipxservice delete <service_type> <service_name> <netid> <node_address> <socket>	Deletes an IPX service from the service table.
configure vlan <name> xnetid <netid> [enet_ii | enet_8023 | enet_8022 | enet_snap]	Configures a VLAN to run IPX routing. Specify: • `enet_ii`—Uses Ethernet 2 header. • `enet_8023`—Uses IEEE 802.3 length field, but does not include the IEEE 802.2 LLC header. • `enet_8022`—Uses IEEE the format and the IEEE 802.2 LLC header. • `enet_snap`—Adds Subnetwork Access Protocol (SNAP) header to IEEE 802.2 LLC header.
enable type20 forwarding {vlan <name>}	Enables the forwarding of IPX type 20 (NetBIOS§ inside IPX) packets from one or more ingress VLANs. The default setting is disabled.
xping {continuous} {size <n>} <netid> <node_address>	Pings an IPX node. If `continuous` is not specified, 4 pings are sent. The default ping packet size is 256 data bytes. The size can be configured to between 1 and 1,484 bytes.

Table 16.4 describes the commands used to configure the IPX route table. For more command options, press the Tab key in the command line interface.

Table 16.4: IPX§ /RIP Configuration Commands

Command	Description
configure ipxrip add vlan [<name> | all]	Configures one or all IPX VLANs to run IPX/RIP. IPX/RIP is enabled by default when you configure the IPX VLAN.

Table 16.4: IPX§ /RIP Configuration Commands (continued)

Command	Description
configure ipxrip vlan [all | <name>] [import-filter | export-filter | trusted-gateway] [none | <access-profile>]	Configures the import, export, or trusted-gateway options and specifies an access profile.
configure ipxrip delete vlan [<name> | all]	Disables IPX/RIP on one or all interfaces.
configure ipxrip vlan [<name> | all] delay <msec>	Configures the time between each IPX/RIP packet within an update interval. The default setting is 55 milliseconds.
configure ipxrip vlan [<name> | all] max-packet-size <size>	Configures the maximum transmission unit (MTU) size of the IPX/RIP packet. The default setting is 432 bytes.
enable ipxrip	Enables IPX/RIP on the router.
configure ipxrip vlan [<name> | all] update-interval <time> {hold-multiplier <number>}	Configures the update interval and hold multiplier for IPX/RIP updates. This setting affects both the periodic update interval of IPX/RIP and the aging interval of learned routes. The default update interval is 60 seconds. The default multiplier is 3.The aging period is calculated using the formula*:* *update-interval x multiplier = aging period*

Table 16.5 describes the commands used to configure IPX/SAP. For more command options, press the Tab key in the command line interface.

Table 16.5: IPX§/SAP Configuration Commands

Command	Description
configure ipxsap add vlan [<name> | all]	Configures an IPX VLAN to run IPX/SAP routing. If no VLAN is specified, all VLANs are configured to run IPX/SAP routing. IPX/SAP routing is enabled by default when the IPX VLAN is configured.

Table 16.5: IPX§/SAP Configuration Commands (continued)

Command	Description
configure ipxsap delete vlan [<name> | all]	Disables IPX/SAP on an interface.
configure ipxsap vlan [<name> | all] delay <msec>	Configures the time between each SAP packet within an update interval. The default setting is 55 milliseconds.
configure ipxsap vlan [<name> | all] [export-filter | import-filter | trusted-gateway] [<access profile> | none]	Configures the import, export, or trusted-gateway options and specifies an access profile.
configure ipxsap vlan [<name> | all] max-packet-size <number>	Configures the MTU size of the IPX/SAP packets. The default setting is 432 bytes.
configure ipxsap vlan [<name> | all] update-interval <time> {hold-multiplier <number>}	Configures the update interval and hold multiplier for IPX/SAP updates. This setting affects both the periodic update interval of SAP and the aging interval of learned routes. The default update interval is 60 seconds.The default multiplier is 3. The aging period is calculated using the formula: *update-interval * multiplier = aging period* Triggered update is always enabled; therefore, new information is processed and propagated immediately.
configure ipxsap vlan <name> gns-delay <msec>	Configures the amount of time the switch waits before answering a GNS request. By default, the switch answers a GNS request as soon as possible (0 milliseconds).
enable ipxsap	Enables IPX/SAP on the router.
enable ipxsap gns-reply {vlan <name>}	Enables GNS-reply on one or all IPX interfaces. If no VLAN is specified, GNS-reply is enabled on all IPX interfaces. The default setting is aging time (in seconds).

IPX Configuration Example

Figure 16.2 builds on the example showing the IP/RIP configuration that was used in Figure 13.4 on page 241. Now, along with having IP VLANs configured, this example illustrates a switch that has two IPX VLANs defined.

The first VLAN is ***Exec*** with these characteristics:

- Protocol-sensitive VLAN using the IPX protocol with the filter IPX_8022
- Ports 4 and 5 have been assigned to ***Exec***
- ***Exec*** is configured for IPX NetID 2516 and IPX encapsulation type 802.2

The second VLAN is ***Support*** with these characteristics:

- Port 7 is assigned to ***Support***
- ***Support*** is configured for IPX NetID A2B5 and IPX encapsulation type 802.2

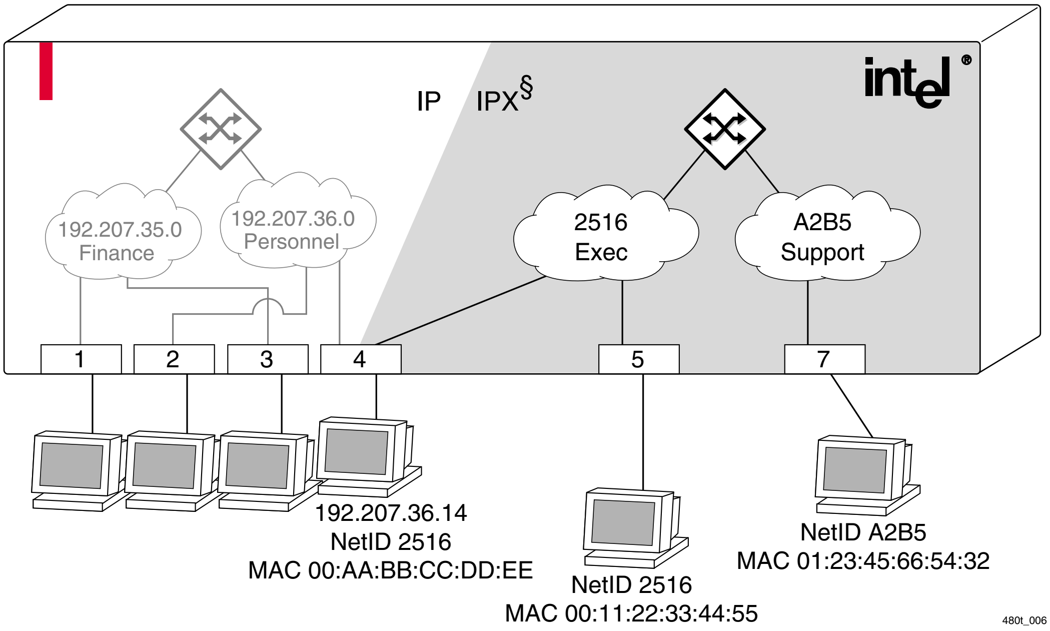

Figure 16.2: IPX routing configuration example

The stations connected to the system generate a combination of IP traffic and IPX traffic. The IP traffic is filtered by the IP VLANs. IPX traffic is filtered by the IPX VLANs.

In this configuration, all IP traffic from stations connected to ports 1 and 3 have access to the IP router through the VLAN ***Finance***. IP traffic on ports 2 and 4 reach the IP router using the VLAN ***Personnel***.

Similarly, IPX traffic from stations connected to ports 4 and 5 have access to the IPX router using the VLAN ***Exec***. IPX traffic to port 7 reaches the IPX router using the VLAN ***Support***. Both ***Exec*** and ***Support*** use enet_8022 as the encapsulation type.

The IPX configuration shown in the example in Figure 16.2 uses these commands:

```
create vlan Exec
create vlan Support
configure Exec protocol ipx_8022
configure Exec add port 4,5
configure Support add port 7
configure Support protocol ipx_8022
configure Exec xnetid 2516 enet_8022
configure Support xnetid A2B5 enet_8022
```

Displaying IPX Settings

To display settings for various IPX components, use the commands listed in Table 16.6. For more command options, press the Tab key in the command line interface.

Table 16.6: IPX§ Show Commands

Command	Description
show ipxconfig {vlan <name>}	Displays IPX configuration information for one or all VLANs.
show ipxfdb	Displays the hardware IPX FDB information.
show ipxrip {vlan <name>}	Displays IPX/RIP configuration and statistics for one or all VLANs.

Table 16.6: IPX§ Show Commands (continued)

Command	Description
show ipxroute {vlan <name> | xnetid <netid> | origin [static | rip | local]}	Displays the IPX routes in the route table.
show ipxsap {vlan <name>} {stats}	Displays IPX/SAP configuration and status for one or all VLANs.
show ipxservice {vlan <name> | name <service name> | type <hex> | origin [static | ipxsap]}	Displays IPX services learned through SAP.
show ipxstats {vlan <name>}	Displays IPX packet statistics for the IPX router, and one or all VLANs.

Resetting and Disabling IPX

To return IPX settings to their defaults and disable IPX functions, use the commands listed in Table 16.7.

Table 16.7: IPX§ Reset and Disable Commands

Command	Description
disable ipxrip	Disables IPX/RIP on the router.
disable ipxsap	Disables IPX/SAP on the router.
disable ipxsap gns-reply {vlan <name>}	Disables GNS reply on one or all IPX interfaces.
disable type20 forwarding {vlan <name>}	Disables the forwarding of IPX type-20 packets.
unconfigure ipxrip {vlan <name>}	Resets the IPX/RIP settings on one or all VLANs to the default. Removes import and export filters, and resets the MTU size, update interval, and inter-packet delay.

Table 16.7: IPX[§] Reset and Disable Commands (continued)

Command	Description
unconfigure ipxsap {vlan <name>}	Resets the IPX/SAP settings on one or all VLANs to the default. Removes import and export filters, and resets the MTU size, update interval, and inter-packet delay.
unconfigure vlan <name> xnetid	Removes the IPX NetID of a VLAN.

17 Access Policies

This chapter describes access policies, and how they are created and implemented on the Intel® NetStructure™ 480T routing switch.

Overview of Access Policies

Access policies are a generalized category of features that impact forwarding and route forwarding decisions. Access policies are used primarily for security and quality of service (QoS) purposes.

There are three categories of access policies:

- IP access lists
- Routing access policies
- Route maps

IP Access Lists

IP access lists consist of IP access rules, and are used to perform packet filtering and forwarding decisions on incoming traffic. They are based on criteria that involves Layer 3 IP or Layer 4 socket source or destination information. Each packet arriving on an ingress port is compared to the access list in sequential order, and is either forwarded to a specified QoS profile or dropped. Using access lists has no impact on switch performance.

Access lists are typically applied to traffic that crosses Layer 3 router boundaries, but it is possible to use access lists within a Layer 2 VLAN.

Routing Access Policies

Routing access policies are used to control the advertisement or recognition of routing protocols, such as Router Information Protocol (RIP), Open Shortest Path First (OSPF) or Border Gateway Protocol (BGP). You can use routing access policies to hide entire networks, or to trust only specific sources for routes or ranges of routes.

The capabilities of routing access policies are specific to the type of routing protocol involved, but are sometimes more efficient and easier to implement than access lists.

IPX§ Routing Access Policies

Routing access policies support IPX, IPX/ RIP, IPX/SAP, and IPX node rules. Routing access policies consist of access rules, and are used to perform packet filtering and forwarding decisions on incoming traffic. Each IPX/RIP or IPX/SAP packet arriving on an ingress port is compared to each access profile rule in sequence, and is either forwarded or dropped. To create IPX access profiles, use this command:

```
create access-profile <access_profile> type
[ipaddress | ipx-node | ipx-net | ipx-sap | as-path |
bgp-community]
```

To configure an IPX net, node or SAP access profile, use this command:

```
configure access-profile <access_profile> [add |
delete] {seq-number} ipx-net <ipx_net_id_in_hex>
<ipx_net_id_mask_in_hex>
```

```
configure access-profile <access_profile> [add |
delete] {seq-number} ipx-node <ipx_net_id_in_hex>
<ipx_net_id_mask_in_hex>
<ipx_node_id_in_mac_address_format>
```

```
configure access-profile <access_profile> [add |
delete] {seq-number} ipx-sap <ipx_sap_type_in_hex>
<ipx_name_string>
```

To assign IPX access profiles as either import or export filters to RIP or SAP, use these commands:

```
configure ipxrip vlan [<vlan name> | all] import-
filter [<access_profile> | none]
```

```
configure ipxrip vlan [<vlan name> | all] export-
filter [<access_profile> | none]
```

```
configure ipxsap vlan [<vlan name> | all] import-
filter [<access_profile> | none]
```

```
configure ipxsap vlan [<vlan name> | all] export-
filter [<access_profile> | none]
```

To view your access profile configuration, use this command:

```
show access-profile <access_profile>
```

Route Maps

Route maps are used to modify or filter routes redistributed into BGP. They are also used to modify or filter the routing information exchanged with BGP neighbors.

Using IP Access Lists

Each entry that makes up the IP access list of the 480T routing switch contains a unique name. It can contain a unique precedence number, as well. Precedence numbers range from 1 to 25,600, with the number 1 having the highest precedence.

The precedence number determines the order in which each criteria rule is examined by the switch. Once a matching entry in the access list is found, the packet is acted on and either forwarded or dropped.

The rules of an IP access list consist of a combination of these six components:

- IP source address and mask
- IP destination address and mask
- TCP or UDP source port range
- TCP or UDP destination port range

- Physical source port
- Precedence number (optional)

How IP Access Lists Work

For each access list entry, you can either permit the packet to be forwarded, or deny the packet (in which case, it is dropped). When you create a permit access list condition, you can optionally specify a QoS profile.

The QoS profile informs the 480T routing switch which bandwidth management and priority to use when transmitting the packet.

When a packet arrives on an ingress port, the packet is compared with the access list rules to determine a match. When a match is found, the packet is processed.

If the access list is of type deny, the packet is dropped. If the list is of type permit, the packet is forwarded. A permit access list can also apply a QoS profile to the packet.

Precedence Numbers

The precedence number is optional, and determines the order in which each rule is examined by the 480T routing switch. Access list entries that contain a precedence number are evaluated from highest to lowest precedence.

You can specify overlapping rules; however, if you are using precedence numbers, overlapping rules without precedence numbers are ignored. Therefore, precedence numbers must be specified among all overlapping rules.

If a new rule without a precedence number is entered, and it overlaps existing rules, the switch rejects the new rule and resolves the precedences among all remaining overlapping rules.

Specifying a Default Rule

To begin constructing an access list, you should specify a default rule. A default rule contains wildcards for destination and source IP address, with no Layer 4 information.

A default rule determines whether the behavior of the access list is an implicit deny or implicit accept. If no access list entry is satisfied, the default rule is used to determine whether the packet is forwarded or dropped. If no default rule is specified, the default implicit behavior is to forward the packet.

This example shows a default entry used to specify an implicit deny:

```
create access-list denyall ip destination 0.0.0.0/0
source 0.0.0.0/0 deny ports any
```

Once the default behavior of the access list is established, you can create additional entries with precedence. The optional precedence numbers range from 1 to 25,600 (number 1 having the highest precedence).

The access list example below performs packet filtering in the following order, as determined by the precedence number:

1. Deny UDP port 16 and TCP port 15 traffic to the 10.2.X.X network.
2. All other TCP port 15 traffic destined for other 10.X.X.X networks is permitted using QoS profile Qp4.
3. All remaining traffic to 10.2.0.0 uses QoS profile Qp3.

With no default rule specified, all remaining traffic is allowed using the default QoS profile.

```
create access-list deny102_16 udp dest 10.2.0.0/8
ip-port 16 source any ip-port any deny ports any
precedence 10

create access-list deny102_15 tcp dest 10.2.0.0/8
ip-port 15 source any ip-port any deny ports any
precedence 20

create access-list allow10_15 tcp dest 10.0.0.0/8
ip-port 15 source any ip-port any permit
qosprofile qp4 ports any precedence 30

create access-list allow102 ip dest 10.2.0.0/8
source 0.0.0.0/0 permit qosprofile qp3 ports any
precedence 40
```

The Permit-Established Keyword

Access lists support the use of the `permit-established` keyword. This keyword allows directional control of attempts to open a TCP session. You can explicitly permit or block session initiation using the keyword. For example, you could use this entry to permit TCP sessions originated from anywhere in the 10.1.0.0 network only:

```
create access-list TCPout tcp destination 10.1.0.0/
16 ip-port any source 0.0.0.0/0 ip-port any
permit-established ports any
```

In this example, using the `permit-established` keyword allows only TCP packets with the ACK (acknowledgement) or RST (reset) bit set to destination 10.1.0.0. from anywhere, but not to any other destination.

Adding and Deleting Access List Entries

You can add and delete entries in the access list. To add an entry, you must supply a unique name and, optionally, a unique precedence number.

To modify an existing entry, you must delete the entry and retype it, or create a new entry with a new unique name.

To delete an access list entry, use the command:

```
delete access-list <name>
```

Maximum Entries

You can use up to 255 entries with an assigned precedence. Along with the 255 entries, you can also create entries that do not use precedence, with these restrictions:

- A source IP address must use wildcards or be completely specified (32-bit mask).
- The Layer 4 source and destination ports must use wildcards or be completely specified (no ranges).
- No physical source port can be specified.

Access Lists for ICMP

Access lists for ICMP (Internet Control Message Protocol) traffic processing are handled somewhat differently. An access list for ICMP is only effective for traffic routed by the switch.

ICMP traffic can either be forwarded (routed) by the switch or discarded, but cannot contain options for assigning a QoS profile.

Other included configuration options for filtering ICMP include:

- IP source and destination address and mask
- ICMP type code
- Physical source port (optional)
- Numbered precedence (optional)

When using an access control list with an IP deny any rule, all ICMP traffic will not be blocked (for either Layer 2 or Layer 3). To block all traffic within Layer 2 and Layer 3, two access lists must be created, an IP deny any rule and an ICMP deny any rule.

Security and Access Policies

ICMP ACL Precedence You can assign precedence values to access lists for ICMP traffic. The precedence number is optional; access list entries that contain a precedence number are evaluated from highest to lowest precedence. Precedence numbers range from 1 to 25,600, with the number 1 having the highest precedence. Assigning precedence allows the switch to resolve conflicts between ICMP rules.

ICMP Deny Rule If an ICMP deny rule is created with type configured as zero, all ICMP traffic with any other type is blocked. The ICMP type zero and code zero is treated as a wildcard and will apply to all ICMP rules.

Verifying Access List Configurations

To verify access list settings you can view the access list configuration to see real-time statistics where access list entries are affected when processing traffic. To view the access list configuration and statistics screen, use this command:

```
show access-list {name | port <port>}
```

To refresh the access list statistics display, use this command:

```
show access-list-monitor
```

Access List Commands

Table 17.1 describes the commands used to configure IP access lists. For further command options, press the Tab key in the command line interface.

Table 17.1: Access List Configuration Commands

Command	Description
create access-list <name> ip destination [<dst_ipaddress>/<dst_mask> | any] source [<src_ipaddress>/<src_mask> | any] [deny | permit <qosprofile> | deny] ports [<portlist> | any] {precedence <number>}	Creates a named IP access list. The access list is applied to all ingress packets. Options include: • `<name>`—Specifies the access list name. The access list name can be between 1 and 16 characters. • `ip`—Specifies an IP access list. • `destination`—Specifies an IP destination address and subnet mask. A mask length of 32 indicates a host entry. An IP address of 0.0.0.0 is a wildcard and matches all. • `source`—Specifies an IP source address and subnet mask. An IP address of 0.0.0.0 is a wildcard and matches all. • `permit`—Specifies that the packets matching the access list description are permitted to be forwarded by this switch. An optional Quality of Service (QoS) profile can be assigned to the access list, to enable the switch to prioritize packets accordingly. • `deny`—Specifies that the packets matching the access list description are filtered (dropped) by the switch. • `precedence`—Specifies the access list precedence number. The range is 1 to 25,600.

Table 17.1: Access List Configuration Commands (continued)

Command	Description
create access-list <name> tcp destination [<dst_ipaddress>/<dst_mask> | any] ip-port [<dst_port> | range <dst_port_min> <dst_port_max> | any] source [<src_ipaddress>/<src_mask> | any] ip-port [<src_port> | range <src_port_min> <src_port_max> | any] [permit <qosprofile> | permit-established | deny] ports [<portlist> | any] {precedence <precedence_num>} {log}	Creates a named IP access list to look at TCP port numbers. The access list is applied to all ingress packets. Options include: • `<name>`—Specifies the access list name. The access list name can be between 1 and 16 characters. • `tcp`—Specifies an IP access list that looks at TCP port numbers. • `destination`—Specifies an IP destination address and subnet mask. A mask length of 32 indicates a host entry. An IP address of 0.0.0.0 is a wildcard and matches all. • `source`—Specifies an IP source address and subnet mask. An IP address of 0.0.0.0 is a wildcard and matches all. • `permit-established`—Specifies that a uni-directional session establishment is allowed. • `permit`—Specifies that the packets matching the access list description are permitted to be forwarded by this switch. An optional QoS profile can be assigned to the access list, to enable the switch to prioritize packets accordingly. • `range`—Specifies the TCP or UDP port range. • `deny`–Specifies that the packets matching the access list description are filtered (dropped) by the switch. • `precedence`—Specifies the access list precedence number. The range is 1 to 25,600.

Table 17.1: Access List Configuration Commands (continued)

<table>
<tr><th>Command</th><th>Description</th></tr>
<tr><td>create access-list <name> udp destination [<dst_ipaddress>/<dst_mask> | any] ip-port [<dst_port> | range <dst_port_min> <dst_port_max> | any] source [<src_ipaddress>/<src_mask> | any] ip-port [<src_port> | range <src_port_min> <src_port_max> | any] [permit <qosprofile> | deny] ports [<portlist> | any] {precedence <precedence_num>}</td><td>Creates a named IP access list to look at UDP port numbers. The access list is applied to all ingress packets. Options include:
• <name>—Specifies the access list name. The access list name can be between 1 and 16 characters.
• udp—Specifies an IP access list that looks at UDP port numbers.
• destination—Specifies an IP destination address and subnet mask. A mask length of 32 indicates a host entry.
• source—Specifies an IP source address and subnet mask.
• permit—Specifies that the packets matching the access list description are permitted to be forward by this switch. An optional QoS profile can be assigned to the access list,to enable the switch to prioritize packets accordingly.
• range—Specifies the TCP or UDP port range.
• deny–Specifies that the packets matching the access list description are filtered (dropped) by the switch.
• precedence—Specifies the access list precedence number. The range is 1 to 25,600.</td></tr>
</table>

Table 17.1: Access List Configuration Commands (continued)

Command	Description
create access-list icmp destination [<dest_ipaddress>/<mask> | any] source [<src_ipaddress>/<source_mask> | any] type <icmp_type> code <icmp_code> [permit | deny] {<portlist>} {precedence <number>}	Creates a named ICMP access list. The access list is applied to all ingress packets. Options include: • `<name>`—Specifies the access list name of between 1 and 16 characters. • `icmp`—Specifies an ICMP access list. • `destination`—Specifies an IP destination address and subnet mask. A mask length of 32 indicates a host entry. • `source`—Specifies an IP source address and subnet mask. • `type`—Specifies the ICMP_TYPE number from 0 to 255. • `code`—Specifies the ICMP_CODE number from 0 to 255. • `permit`—Specifies that packets matching the access list description are forwarded. An optional QoS profile can be assigned to the access list, so the switch can prioritize packets accordingly. • `deny`—Specifies that packets matching the access list description are filtered (dropped) by the switch.
delete access-list <name>	Deletes an access list.
disable access-list <name> [counter | log]	Disables the collection of access-list statistics.
enable access-list <name> [counter | log]	Enables the collection of access-list statistics. The default setting is enabled.
disable access-list <name> log	Disables logging of a message (with details of packet properties) to the Syslog facility for each packet that matches the access list description.
enable access-list <name> log	Enables logging of message, (with details of packet properties) to the Syslog facility for each packet matching the access list description.

Table 17.1: Access List Configuration Commands (continued)

Command	Description
show access-list {<name> | ports <portlist>}	Displays access-list information.
show access-list-fdb	Displays the hardware access control list mapping.
show access-list-monitor	Refreshes the access-list statistics display.

IP Access List Examples

This section presents two IP access list examples:

- Using the permit-establish keyword
- Filtering ICMP packets

Example 1: Using the Permit-Established Keyword

This example uses an access list that permits TCP sessions (Telnet, FTP, and HTTP) to be established in one direction.

The switch shown in Figure 17.1 is configured as:

- Two VLANs, NET10 VLAN and NET20 VLAN, are defined.
- The IP address for NET10 VLAN is 10.10.10.1/24.
- The IP address for NET20 VLAN is 10.10.20.1/24.
- The workstations are configured using addresses 10.10.10.100 and 10.10.20.100.
- IP Forwarding is enabled.

These sections detail the steps used to configure the example.

Step 1 – Deny IP Traffic

First, create an access-list that blocks all IP-related traffic. This includes any TCP- and UDP-based traffic. Although ICMP is used in conjunction with IP, it is technically not an IP data packet. Thus, ICMP data traffic, such as ping traffic, is not affected.

Use this command to create the access-list:

```
create access-list denyall ip destination any
source any deny ports any
```

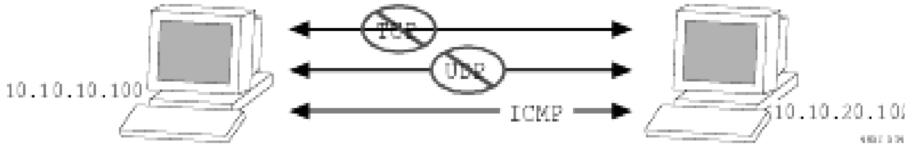

Figure 17.1: Access list denies all TCP and UDP traffic

Step 2 – Allow TCP Traffic

The next set of access-list commands permits TCP-based traffic to flow. Because each session is bidirectional, an access-list must be defined for each direction of the traffic flow. UDP traffic is still blocked.

Use these commands to create the access list defined for bidirectional traffic flow:

```
create access-list tcp1 tcp destination 10.10.20.100/
32 ip any source 10.10.10.100/32 ip any permit qp1
ports any precedence 20
```

```
create access-list tcp2 tcp destination
10.10.10.100/32 ip any source 10.10.20.100/32 ip
any permit qp1 ports any precedence 21
```

Figure 17.2 illustrates the outcome of this access list.

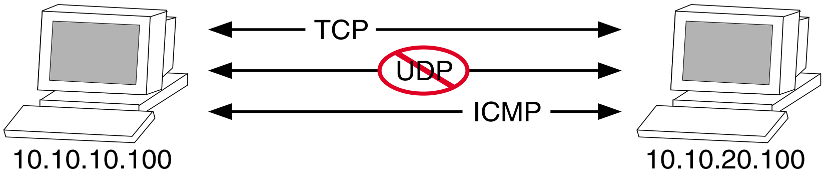

Figure 17.2: Access list allows TCP traffic

Step 3 - Permit-Established Access List

When a TCP session begins, there is a three-way handshake that includes a sequence of a SYN, SYN/ACK and ACK packets. Figure 17.3 shows an illustration of the handshake that occurs when Host A initiates a TCP session to Host B. After this sequence, actual data can be passed.

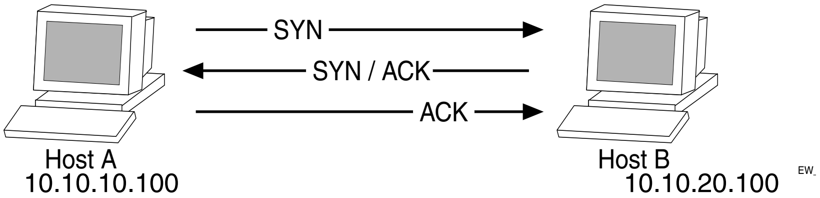

Figure 17.3: Host A initiates a TCP session to Host B

An access list that uses the permit-established keyword filters the SYN packet in one direction.

Use the permit-established keyword to allow only Host A to be able to establish a TCP session to Host B and to prevent any TCP sessions from being initiated by Host B, as illustrated in Figure 17.3. The syntax for this access-list is:

Pay attention to the destination and source address, and the desired effect.

```
create access-list <mylist> tcp destination
<ipaddress> ip-port <portnumber> source <ipaddress>
ip-port any permit-established ports <portnumber>
precedence 8
```

The exact command line entry for this example is:

This rule has a higher precedence than the rule tcp2*.*

```
create access-list telnet-allow tcp destination
10.10.10.100/32 ip-port 15 source any ip-port any
permit-established ports any precedence 8
```

Figure 17.4 shows the final outcome of this access list.

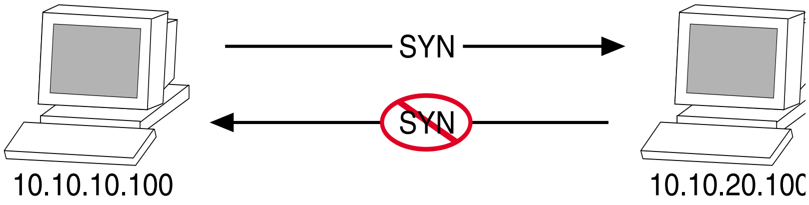

Figure 17.4: Permit-established access list filters out SYN packet to destination

Example 2: Filtering ICMP Packets

This example creates an access list that filters out ping (ICMP echo) packets. ICMP echo packets are defined as type **`any`** code **`any`**.

The command to create this access list is:

```
create access-list denyping icmp destination any
source any type any code any deny ports any
```

Figure 17.5 shows the final outcome of this access list.

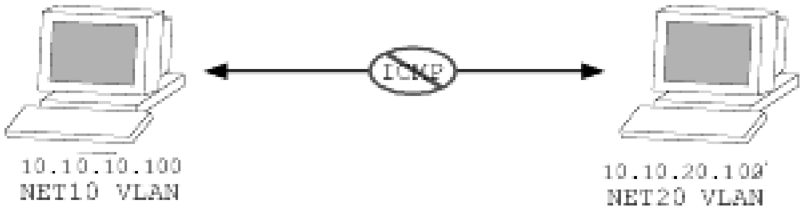

Figure 17.5: ICMP packets are filtered out

Using Routing Access Policies

Access policy entries can be one of these types:

- IP addresses and subnet masks
- VLANs
- Autonomous system path expressions (AS-Path), Border Gateway Protocol (BGP) only
- BGP communities (BGP only)

See “Creating an Access Profile” on page 324.

To use routing access policies

1. Create an access profile.
2. Configure the access profile to be of type permit, deny, or none.
3. Add entries to the access profile.
4. Apply the access profile.

Creating an Access Profile

The first thing to do when using routing access policies is to create an access profile. An access profile has a unique name, and contains one of these entry types:

- A list of IP addresses and associated subnet masks
- One or more autonomous system path expressions (BGP only)
- One or more BGP community numbers (BGP only)

You must give the access profile a unique name (in the same manner as naming a VLAN, protocol filter, or Spanning Tree Domain). You must also indicate the type of access list.

To create an access profile, use this command:

```
create access-profile <access_profile> type
[ipaddress | as-path | bgp-community]
```

Configuring an Access Profile Mode

After the access profile is created, you must configure the access profile mode. The access profile mode determines whether the items in the list are to be permitted access or denied access.

There are three available modes:

- Permit—The **`permit`** access profile mode permits the operation, if it matches any entry in the access profile. If the operation does not match any entries in the list, the operation is denied.
- Deny—The **`deny`** access profile mode denies the operation, if it matches any entry in the access profile. If it does not match all specified entries in the list, the operation is permitted.
- None—Using the none mode, the access profile can contain a combination of **`permit`** and **`deny`** entries. Each entry must include a **`permit`** or **`deny`** attribute. The operation is compared

with each entry in the list. Once a match is found, the operation is either permitted or denied, depending on the configuration of the matched entry. If no match is found, the operation is implicitly denied.

To add or delete IP addresses or VLANs from an access profile, use this command:

```
configure access-profile <access_profile> [add |
delete] {ipaddress <ipaddress> <mask>}
```

Then, configure the access profile mode using

```
configure access-profile <access_profile> mode
[permit | deny | none]
```

Adding an Access Profile Entry

Next, configure the access profile by adding or deleting IP addresses, autonomous system path expressions, or BGP communities, using this command:

```
configure access-profile <access_profile> [add |
delete | mode] {<seq_number>} {permit | deny}
[ipaddress <ipaddress> | <mask> {exact} | as-path
<path-expression> | bgp-community [internet | no-
export | no-advertise | no-export-subconfed |
<as_no:number> | number <community>]]
```

These sections describe the **`configure access-profile add`** command.

Specifying Subnet Masks

The subnet mask specified in the access profile command is interpreted as a reverse mask. A reverse mask indicates the bits that are significant in the IP address. In other words, a reverse mask specifies the part of the address that must match the IP address to which the profile is applied.

If you configure an IP address that is an exact match, specifically denied or permitted, use a mask of /32 (for example, 141.251.24.28/32).

If the IP address represents all addresses in a subnet address that you wish to **`deny`** or **`permit`**, then configure the mask to cover only the subnet portion (for example, 141.251.10.0/24). The keyword **`exact`** can be used when you wish to match only against the subnet address, and ignore all addresses within the subnet.

If you are using CIDR subnet masking, the same logic applies, but the configuration is more tricky. For example, the address 141.251.24.128/25 represents any host from network 141.251.24.128/255.255.255.128.

Sequence Numbering

You can specify the sequence number for each access profile entry. If you do not specify a sequence number, entries are sequenced in the order they are added. Each entry is assigned a value of five more than the sequence number of the last entry.

Permit and Deny Entries

If you have configured the access profile mode to be **`none`**, you must specify each entry type as either **`permit`** or **`deny`**. If you do not specify the entry type, it is added as a **`permit`** entry. If you have configured the access profile mode to be **`permit`** or **`deny`**, it is not necessary to specify a type for each entry.

Autonomous System Expressions

The **`as-path`** keyword uses a regular expression string to match against the AS-path. Regular expression notation can include any of the characters listed in Table 17.2.

Table 17.2: Regular Expression Notation

Character	Definition
[,]	Specifies a range of numbers to be matched
.	Matches any number
^	Matches the beginning of the AS path
$	Matches the end of the AS path
–	Matches the beginning or end, or a space
-	Separates the beginning and end of a range of numbers

Table 17.2: Regular Expression Notation

Character	Definition
*	Matches zero or more instances
+	Matches one or more instances
?	Matches zero or one instance

Deleting an Access Profile Entry

To delete an access profile entry, use this command:

```
configure access-profile <access_profile> delete
<seq_number>
```

Applying Access Profiles

After the access profile is defined, apply it to one or more routing protocols or VLANs. When an access profile is applied to a protocol function (for example, the export of RIP routes) or a VLAN, this forms an access policy.

A profile can be used by multiple routing protocol functions or VLANs, but a protocol function or VLAN can use only one access profile.

Routing Access Policies for RIP

If the RIP protocol is being used, you can configure the 480T routing switch to use an access profile to determine any of these:

- **Trusted Neighbor**—Use an access profile to determine trusted RIP router neighbors for the VLAN on the switch running RIP. To configure a trusted neighbor policy, use this command:

  ```
  configure rip vlan [<name> | all] trusted-
  gateway [<access_profile> | none]
  ```

- **Import Filter**—Use an access profile to determine which RIP routes are accepted as valid routes. You can combine this policy with the trusted neighbor policy to accept selected routes only from a set of trusted neighbors. To configure an import filter policy, use this command:

```
configure rip vlan [<name> | all] import-filter
[<access_profile> | none]
```

- **Export Filter**—Use an access profile to determine which RIP routes are advertised into a particular VLAN, using this command:

```
configure rip vlan [<name> | all] export-filter
[<access_profile> | none]
```

Examples

In the example shown in Figure 17.6, a switch is configured with three VLANs, ***Engsvrs, Sales*** and ***Backbone***. The RIP protocol is used to communicate with other routers on the network. The administrator wants to allow internal access to all the VLANs on the switch, but no access to the router that connects to the Internet. The remote router that connects to the Internet has a local interface connected to the corporate backbone. The IP address of the local interface connected to the corporate backbone is 10.0.0.10/24.

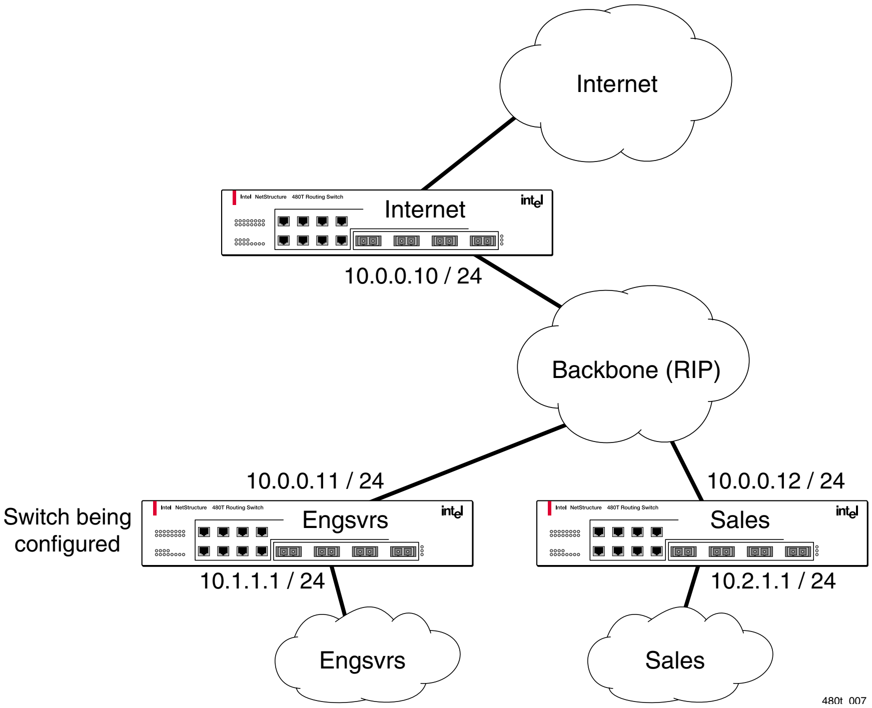

Figure 17.6: RIP access policy example

Assuming the backbone VLAN interconnects all the routers in the company (and, therefore, the Internet router does not have the best routes for other local subnets), the commands to build the access policy for the switch would be:

```
create access-profile nointernet type ipaddress

configure access-profile nointernet mode deny

configure access-profile nointernet add ipaddress
10.0.0.10/32

configure rip vlan backbone trusted-gateway
nointernet
```

If the administrator wants to restrict any user belonging to the VLAN ***Engsvrs*** from reaching the VLAN ***Sales*** (IP address 10.2.1.0/24), the additional access policy commands to build the access policy would be:

```
create access-profile nosales type ipaddress

configure access-profile nosales mode deny

configure access-profile nosales add ipaddress
10.2.1.0/24

configure rip vlan backbone import-filter nosales
```

This configuration results in the switch having no route back to the VLAN ***Sales***.

Routing Access Policies for OSPF

For information on converting an OSPF area into an IP type format see "OSPF (Open Shortest Path First)" on page 443.

Because OSPF is a link-state protocol, the access policies associated with OSPF are different in nature than those associated with RIP. Access policies for OSPF are intended to extend the filtering and security capabilities of OSPF (for example, link authentication and the use of IP address ranges). If the OSPF protocol is being used, you can configure the switch to use an access profile to determine any of these:

- **Inter-area Filter**—For switches configured to support multiple OSPF areas (an ABR function), you can apply an access profile to an OSPF area that filters a set of OSPF inter-area routes from being sourced from any other areas. To configure an inter-area filter policy, use this command:

  ```
  configure ospf area <area_id> interarea-filter
  [<access_profile> | none]
  ```

- **External Filter**—For switches configured to support multiple OSPF areas (an ABR function), you can apply an access profile to an OSPF area that filters a set of OSPF external routes from being advertised into that area. To configure an external filter policy, use this command:

  ```
  configure ospf area <area_id> external-filter
  [<access_profile> | none]
  ```

If any of the external routes specified in the filter have already been advertised, those routes will remain until the associated LSAs in that area time-out.

- **ASBR Filter**—For switches configured to support RIP and static route re-distribution into OSPF, you can use an access profile to limit the routes advertised into OSPF for the switch as a whole. To configure an ASBR filter policy, use this command:

  ```
  configure ospf asbr-filter [<access_profile> |
  none]
  ```

- **Direct Filter**—For switches configured to support direct route re-distribution into OSPF, you can use an access profile to limit the routes that are advertised into OSPF for the switch as a whole. To configure a direct filter policy, use this command:

  ```
  configure ospf direct-filter [<access_profile> |
  none]
  ```

OSPF Access Policy Example

Figure 17.7 illustrates an OSPF network that resembles the network used previously in the RIP example. In this example, access to the Internet is accomplished using the ASBR function on the switch labeled ***Internet***. As a result, all routes to the Internet are done through external routes.

Suppose the network administrator wishes to only allow access to certain Internet addresses falling within the range 192.1.1.0/24 to get to and from the internal backbone.

To configure the switch labeled ***Internet***, the commands would be:

```
create access-profile okinternet ipaddress
configure access-profile okinternet mode permit
configure access-profile okinternet add 192.1.1.0/
24
configure ospf asbr-filter okinternet
```

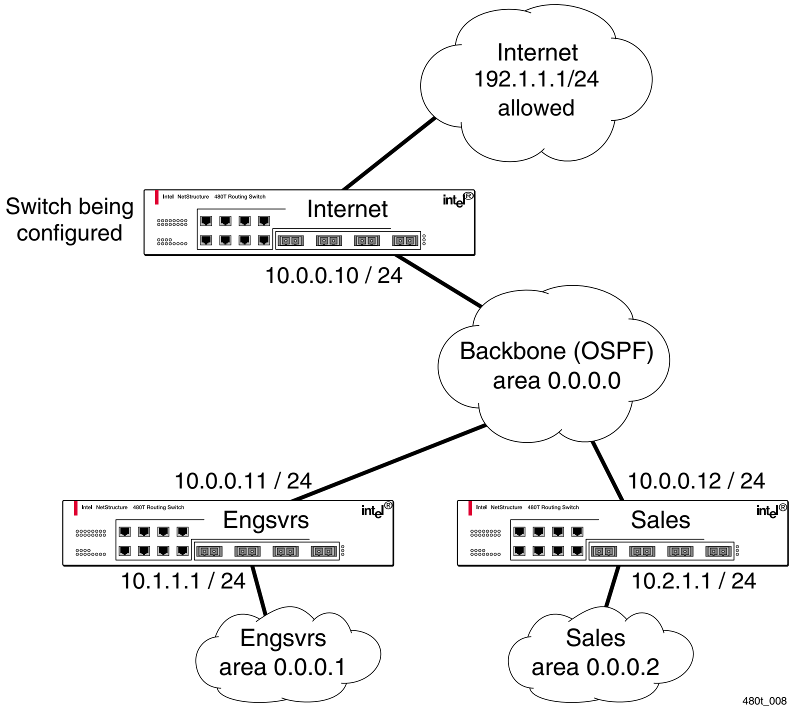

Figure 17.7: OSPF access policy example

Routing Access Policies for DVMRP

The access policy capabilities for DVMRP resemble those for RIP. If the DVMRP protocol is used for routing IP multicast traffic, you can configure the switch to use an access profile to determine:

- **Trusted Neighbor**—Use an access profile to determine trusted DVMRP router neighbors for the VLAN on the switch running DVMRP. To configure a trusted neighbor policy, use this command:

```
configure dvmrp vlan [<name> | all] trusted-
gateway [<access_profile> | none]
```

- **Import Filter**—Use an access profile to determine which DVMRP routes are accepted as valid routes. To configure an import filter policy, use this command:

```
configure dvmrp vlan [<name> | all] import-
filter [<access_profile> | none]
```

- **Export Filter**—Use an access profile to determine which DVMRP routes are advertised into a particular VLAN, using this command:

  ```
  configure dvmrp vlan [<name> | all] export-
  filter [<access_profile> | none]
  ```

DVMRP Example

In this example, the network used in the previous RIP example is configured to run DVMRP. The network administrator wants to disallow Internet access for multicast traffic to users on the VLAN ***Engsvrs***.

This is accomplished by preventing the learning of routes that originate from the switch labeled ***Internet*** through DVMRP on the switch labeled ***Engsvrs***.

To configure the switch labeled ***Engsvrs***, use these commands:

```
create access-profile nointernet type ipaddress

configure access-profile nointernet mode deny

configure access-profile nointernet add ipaddress
10.0.0.10/32

configure dvmrp vlan backbone trusted-gateway
nointernet
```

Suppose the administrator wants to preclude users on the VLAN ***Engsvrs*** from seeing any multicast streams that are generated by the VLAN ***Sales*** across the backbone. The commands for the additional configuration of the switch labeled ***Engsvrs*** are:

```
create access-profile nosales type ipaddress

configure access-profile nosales mode deny

configure access-profile nosales add ipaddress
10.2.1.0/24

configure dvmrp vlan backbone import-filter nosales
```

Routing Access Policies for PIM

PIM (Protocol Independent Multicasting) leverages the unicast routing capability that is already present in the 480T routing switch. If the PIM protocol is used for routing IP multicast traffic, you can configure the switch to use an access profile to determine trusted neighbor (PIM) router neighbors for the VLAN on the switch running PIM.

To configure a trusted neighbor policy, use this command:

```
configure pim vlan [<name> | all] trusted-
gateway [<access_profile> | none]
```

PIM Example

With PIM, you can use the unicast access policies to restrict multicast traffic. In this example, a network similar to the example used in the previous RIP example is also running PIM. The network administrator wants to disallow Internet access for multicast traffic to users on the VLAN ***Engsvrs***. This is accomplished by preventing the learning of routes that originate from the switch labeled ***Internet*** using PIM on the switch labeled ***Engsvrs***.

To configure the switch labeled ***Engsvrs***, the commands would be:

```
create access-profile nointernet type ipaddress

configure access-profile nointernet mode deny

configure access-profile nointernet add ipaddress
10.0.0.10/32

configure pim vlan backbone trusted-gateway
nointernet
```

Routing Access Policies for BGP

If the BGP protocol is being used, you can configure the switch to use an access profile to determine:

- **NLRI filter**—Use an access profile to determine the NLRI information that must be exchanged with a neighbor. To configure an NLRI filter policy, use this command:

  ```
  configure bgp neighbor [<ipaddress> | all]
  nlri-filter [in | out] [<access_profile> |
  none]
  ```

  The NLRI filter access policy can be applied to the ingress or egress updates, using the `in` and `out` keywords, respectively.

- **Autonomous system path filter—**Use an access profile to determine which NLRI information must be exchanged with a neighbor based on the AS path information present in the path attributes of the NLRI. To configure an autonomous system path filter policy, use this command:

```
configure bgp neighbor [<ipaddress> | all] as-
path-filter [in | out] [<access_profile> |
none]
```

You can apply the autonomous system path filter to the ingress or egress updates, using the **in** and **out** keywords, respectively.

Making Changes to a Routing Access Policy

Changes to profiles applied to OSPF require rebooting the switch or disabling and re-enabling OSPF.

You can change the routing access policy by changing the associated access profile. However, the propagation of the change depends on the protocol and policy involved. Propagation of changes applied to RIP, DVMRP, and PIM access policies depend on the respective protocol timers to age-out entries.

In BGP, the change to the policy is immediately effective on the routing information exchanged after the policy changes. You can apply the changes on the routing information that had been exchanged before the policy changes, by issuing a soft reset on the ingress or egress side, depending on the change.

For soft resets to be applied on the ingress side, the changes must have been previously enabled on the neighbor.

Removing a Routing Access Policy

To remove a routing access policy, you must remove the access profile from the routing protocol or VLAN. All the commands that apply an access profile to form an access policy also have the option of choosing **none** as the access profile. Using the **none** option removes any access profile of that particular type from the protocol or VLAN, and, therefore, removes the access policy.

Routing Access Policy Commands

Table 17.3 describes the commands used to configure routing access policies. Press the Tab key in the command line interface for further command options.

Table 17.3: Routing Access Policy Configuration Commands

Command	Description
configure access-profile <access_profile> add {<seq-number>} {permit | deny} [ipaddress <ipaddress> <mask> {exact} | as-path <path_expression> | bgp-community [internet | no-advertise | no-export | no-export-subconfed | <as_no:number> | number <community>]]	Adds an entry to the access profile. The explicit sequence number, and permit or deny attribute should be specified if the access profile mode is none. Specify: • `<seq-number>`—The order of the entry within the access profile. If no sequence number is specified, the new entry is added to the end of the access profile and is automatically assigned a value of 5 more than the sequence number of the last entry. • `permit	deny`—Per-entry permit or deny specification. The per-entry attribute only takes effect if the access profile mode is `none`. Otherwise, the overall access profile type takes precedence. • `<ipaddress> <mask>`—An IP address and mask. If the attribute `exact` is specified for an entry, then an exact match with address and mask is performed; subnets within the address range do not match entry against entry. • `as-path`—A regular expression string to compare with the autonomous system path. • `bgp-community`—The BGP community number in as_no:number format, or as an unsigned 32-bit integer in decimal format. The BGP community `internet` matches against all routes, because all routes belong to the Internet community.

Table 17.3: Routing Access Policy Configuration Commands (continued)

Command	Description
configure access-profile <access_profile> delete <seq_number>	Deletes an access profile entry using the sequence number.
configure access-profile <access_profile> mode [permit | deny | none]	Configures the access profile to one of the following: • `permit`—Allows the addresses that match the access profile description. • `deny`—Denies the addresses that match the access profile description. • `none`—Permits and denies access on a per-entry basis. Each entry must be added to the profile as either **`permit`** or **`deny`**. The default setting is permit.
configure bgp neighbor [<ipaddress> | all] as-path-filter [in | out] [<access_profile> | none]	Configures BGP to use the AS-path filter for the routing information exchanged with the neighbor.
configure bgp neighbor [<ipaddress> | all] nlri-filter [in | out] [<access_profile> | none]	Configures BGP to use the NLRI filter for routing information exchanged with a neighbor.
configure dvmrp vlan [<name> | all] export-filter [<access_profile> | none]	Configures DVMRP to filter out certain routes while performing the route advertisement.
configure dvmrp vlan [<name> | all] import-filter [<access_profile> | none]	Configures DVMRP to filter certain routes received from its neighbor.
configure dvmrp vlan [<name> | all] trusted-gateway [<access_profile> | none]	Configures DVMRP to use the access policy to determine which DVMRP neighbor is trusted and to receive routes from the access policy.
configure ospf area <area_id> external-filter [<access_profile> | none]	Configures the router to use an access policy or policies to determine which external routes are allowed to be exported into the area. This router must be an ABR.

Table 17.3: Routing Access Policy Configuration Commands (continued)

Command	Description
configure ospf area <area_id> interarea-filter [<access_profile> | none]	Configures the router to use the access policy to determine which inter-area routes are allowed to be exported into the area. This router must be an ABR.
configure ospf asbr-filter [<access_profile> | none]	Configures the router to use the access policy to limit the routes that are advertised into OSPF for the switch as a whole, for switches configured to support RIP and static route redistribution into OSPF.
configure ospf direct-filter [<access_profile> | none]	Configures the router to use the access policy to limit the routes that are advertised into OSPF for the switch as a whole, for switches configured to support direct route redistribution into OSPF.
configure pim vlan [<name> | all] trusted-gateway [<access-profile> | none]	Configures PIM to use the access profile to determine which PIM neighbor is to receive or reject the routes.
configure rip vlan [<name> | all] export-filter [<access-profile> | none]	Configures RIP to suppress certain routes when performing route advertisements.
configure rip vlan [<name> | all] import-filter [<access_profile> | none]	Configures RIP to ignore certain routes received from its neighbor.
configure rip vlan [<name> | all] trusted-gateway [<access_profile> | none]	Configures RIP to use the access list to determine which RIP neighbor is to receive (or reject) the routes.

Using Route Maps

Route maps are a mechanism you can use to conditionally control the redistribution of routes between two routing domains, and to modify the routing information that is redistributed.

Route maps are used in conjunction with the match and set operations. A match operation specifies a criteria that must be matched. A set operation specifies a change that is made to the route when the match operation is successful.

There are three basic steps to configuring a route-map:

1. Create a route-map.

2. Add entries to the route map.

3. Add statements to the route map entries.

Creating a Route Map

To create a route-map , use this command:

```
create route-map <route-map>
```

Add Entries to the Route Map

To add entries to the route map, use this command:

```
configure route-map <route-map> add <sequence number>
[permit | deny] {match-one | match-all}
```

where:

- The unique `sequence number` identifies the entry, and determines the position of the entry in the route map. Route maps are evaluated sequentially.
- The `permit` keyword permits the route; the `deny` keyword denies the route and is applied only if the entry is successful.
- The `match-one` keyword is a logical **or.** The route map is successful if at least one of the matching statements is true.
- The `match-all` keyword is a logical **and.** The route map is successful when all match statements are true. This is the default setting.

Add Statements to the Route Map Entries

To add statements to the route map entries, use one of these three commands:

```
configure route-map <route-map> <sequence number>
add match [nlri-list <access_profile> | as-path
[access_profile <access-profile> | <as num>] |
community [access-profile <access_profile> |
<as_num>:<number> | number <community>] | next-hop
<ipaddress> | med <number> | origin [igp | egp |
incomplete | tag <number>]]
```

```
configure route-map <route-map> <sequence number>
add set [as-path <as_num> | community [remove |
{add | delete} [access-profile <access_profile |
<as_num:number> | number <number>] |] next-hop
<ipaddress> | med <number> | local-preference
<number> | origin [igp | egp | incomplete | tag
<number>]]
```

```
configure route-map <route-map> <sequence number>
add goto <route-map>
```

where:

- The **route-map** is the name of the route map.
- The **sequence number** identifies the entry in the route map to which this statement is being added.
- The **match**, **set**, and **goto** keywords specify the operations to be performed. Within an entry, the statements are sequenced in the order of their operation. The match statements are first, followed by set, and then goto.
- The **nlri-list**, **as-path**, **community**, **next-hop**, **med**, **origin**, and **weight** keywords specify the type of values that must be applied using the specified operation against the corresponding attributes as described in Table 17.4 and Table 17.5.

Table 17.4: Match Operation Keywords

Keyword	Description
nlri-list <access_profile>	Matches the NLRI against the specified access profile.
as-path [<access_profile> | <as-no>]	Matches the AS path in the path attributes against the specified access profile or AS number.

Table 17.4: Match Operation Keywords

Keyword	Description
community [<access_profile> | <community>]	Matches the communities in the path attribute against the specified BGP community access profile or the community number.
next-hop <ipaddress>	Matches the next-hop in the path attribute against the specified IP address.
med <number>	Matches the multi-existing discriminator (MED) in the path attribute against the specified MED number.
origin [igp | egp | incomplete]	Matches the origin in the path attribute against the specified origin.

Table 17.5: Set Operation Keywords

Keyword	Definition
as-path <as no>	Adds the specified AS number to the beginning of the AS path in the path attribute.
community <community>	Adds the specified community to the existing community in the path attribute.
next-hop <ipaddress>	Sets the next hop in the path attribute to a specified IP address.
med <number>	Sets MED in the path attribute to a specified MED number.
local-preference <number>	Sets the local preference in the path attribute to the specified local preference number.
weight <number>	Sets weight associated with NLRI to a specified number.
origin	Sets origin in the path attributes to the specified origin.

Route Map Operation

The entries in the route map are processed in the ascending order of the sequence number. Within the entry, the match statements are processed first. When the match operation is successful, the set and goto statements within the entry are processed, and the action associated with the entry is either applied, or the next entry is processed. If the end of the route map is reached, it is implicitly denied.

When there are multiple match statements, the primitive **`match-one`** or **`match-all`** in the entry determines how many matches are required for success.

When there are no match statements in an entry, the entry is considered a successful match.

Route Map Example

Figure 17.8 shows a topology using route maps to filter or modify routing information that is exchanged between the neighbors RTA and RTB using BGP.

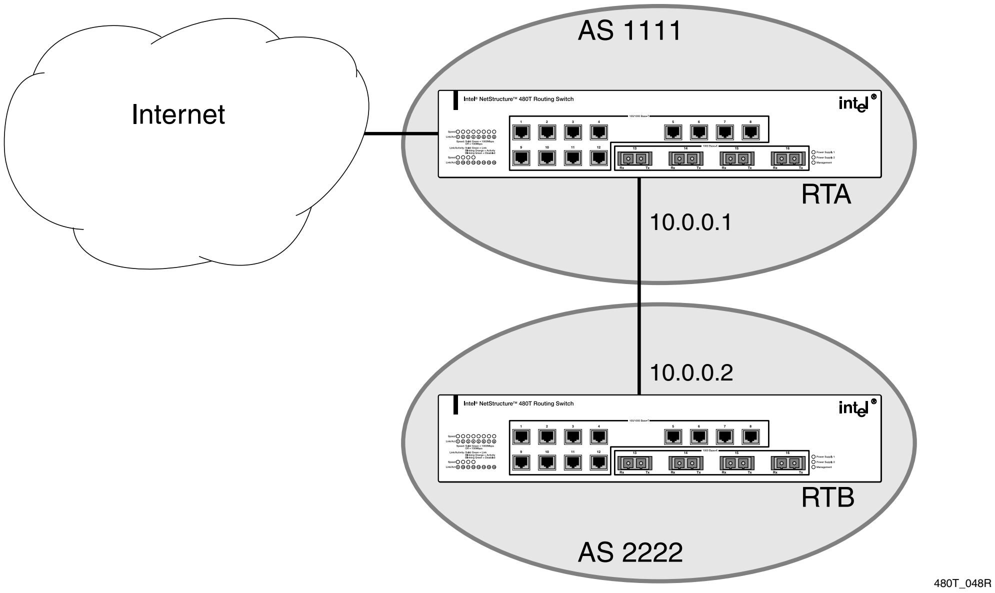

Figure 17.8: Route maps

These points apply to this example:

- RTA is a member of AS 1111 and peers with a router in the Internet to receive the entire Internet routing table.
- RTB is a member of AS 2222, and has an EBGP connection with RTA through which it receives the Internet routing table.
- AS 1111 is acting as a transit AS for all traffic between AS 2222 and the Internet.

  If you are the administrator of AS 1111 and you want to filter out route information about network 221.1.1.0/24 and its subnets from being passed on to AS 2222, you can configure a route-map on the egress side of RTA's EBGP connection with RTB, and filter out the routes.

To configure RTA, use these commands:

```
create access-profile iplist type ipaddress
configure iplist add ipaddress 221.1.1.0/24

create route-map bgp-out
configure bgp-out add 10 deny
configure bgp-out 10 add match nlri-list iplist
configure bgp-out add 20 permit

configure bgp neighbor 10.0.0.2 route-map-filter
out bgp-out
configure bgp neighbor 10.0.0.2 soft-reset out
```

To modify the routing information originated from AS 300 to include a MED value of 200, the sequence of commands would be:

```
create access-profile aslist type as-path
configure aslist add as-path "^300"

configure bgp-out add 15 permit
configure bgp-out 15 add match as-path access-
profile aslist
configure bgp-out 15 add set med 200

configure bgp neighbor 10.0.0.2 soft-reset out
```

Changes to Route Maps

Changes to the route maps used to modify or filter NLRI information exchanged with neighbors is immediately effective on the routing information exchanged after the policy changes.

You can apply the changes on the NLRI information that had been exchanged before the policy changes, by issuing a soft reset on the ingress or egress side, depending on the changes.

For soft resets to be applied on the ingress side, the changes must be previously enabled on the neighbor.

Changes to the route maps associated with network aggregation or redistribution commands become effective after a maximum interval of 30 seconds. You can immediately apply them by using the soft reconfiguration command.

Route Maps in BGP

Route maps are used in BGP to modify or filter NLRI information exchanged with neighbors. They are also used in NLRI information that originates through network command, aggregation, or redistribution.

Route Map Commands

Table 17.6 describes route map commands. For further command options, press the Tab key in the command line interface.

Table 17.6: Route Map Commands

Command	Description
configure route-map <route-map> [add | delete] <sequence number> [deny | permit] {match-all | match one}	Adds or deletes entries to the route map. Specify: • The `sequence number` uniquely identifies the entry, and determines the position of the entry in the route map. Route maps are evaluated sequentially. • The `permit` keyword permits the route; the `deny` keyword denies the route and is applied only if the entry is successful. • The `match-one` keyword is a logical **or.** The route map is successful if at least one of the matching statements is true. • The `match-all` keyword is a logical **and**. The route map is successful when all match statements are true. The `match all` setting is the default.
configure route-map <route-map> <sequence number> add goto <route-map>	Configures a route-map `goto` statement.
configure route-map <route-map> <sequence number> add match [nlri-list <access_profile> | as-path [access_profile <access_profile> | <as_num>] | community [access-profile <access_profile> | <as_num>:<number | number <community>] | next-hop <ipaddress> | med <number> | origin [igp | egp | incomplete | tag <number>]]	Configures a route-map `match` statement. Specify: • `route-map`—The name of the route map. • `sequence number`—The statement in the route map to which this statement is being added. • `nlri-list`, `as-path`, `community`, `next-hop`, `med`, and `origin`—The type of values that must be applied using the specified operation against the corresponding attributes as described in Table 17.4.

Table 17.6: Route Map Commands (continued)

Command	Description
configure route-map <route-map> <sequence number> add set [accounting index <num> value <num> | as-path <as_num> | community [remove | {add | delete} [access-profile <access_profile> | <as_num:number> | number <number>]] | cost <num> | cost-type [<ase-type-1 | ase-type-2>] | next-hop <ipaddress> | med <number> | local-preference <number> | origin [igp | egp | incomplete] | tag <num> | weight <num>]	Configures a route-map `set` statement. Specify: • `route-map -` The name of the route map. • `sequence number` – The statement in the route map to which this statement is being added. • `as-path`, `community`, `next-hop`, `med`, `local-preference`, and `origin -` Specify the type of values that must be applied using the specified operation against the corresponding attributes as described in Table 17.5.
configure route-map <route-map> <sequence number> delete goto <route-map>	Deletes a route-map `goto` statement.
configure route-map <route-map> <sequence number> delete match [nlri-list <access_profile> | as-path [access-profile <access_profile> | <as_no>] | community [access-profile | no-advertise | no-export | no-export-subconfed | number | <AS-id>] | next-hop <ipaddress> | med <number> | origin [igp | egp | incomplete]]	Deletes a route-map `match` statement.
configure route-map <route-map> <sequence number> delete set [accounting index <num> value <num> | as-path <as_num> | community [remove | {add | delete} [access-profile <access_profile> | <as_num:number> | number <number>]] | cost <num> | cost-type [<ase-type-1 | ase-type-2>] | next-hop <ipaddress> | med <number> | local-preference <number> | origin [igp | egp | incomplete] | tag <num> | weight <num>]	Deletes a route-map `set` statement.

Table 17.6: Route Map Commands (continued)

Command	Description
configure route-map <route-map> add <sequence number> [permit | deny] {match-one | match-all]	Adds a statement to the route map with the specified sequence number and action. The sequence number determines the order of the statement in the route map, and the action specifies the action to be taken on a successful match against the statements in the route map.
configure route-map <route-map> delete <sequence number>	Deletes a statement from the route map.
create route-map <route-map>	Creates a route-map statement.
delete route-map <route_map>	Deletes a route-map statement from the route map.

18 Server Load Balancing (SLB)

Overview

The Server Load Balancing (SLB) feature of the Intel® NetStructure™ 480T routing switch divides many client requests among several servers. This activity is transparent to the client using the resource. It is mainly used for Web hosting where several redundant servers are used to increase the performance and reliability of busy Web sites.

Using SLB, the switch can manage and balance traffic for client equipment such as Web servers, cache servers, routers, firewalls, and proxy servers. SLB offers a variety of useful features that meet the special needs of e-commerce sites, Internet service providers, and managers of large intranets.

SLB Components

There are three components that comprise an SLB system:

- Nodes
- Pools
- Virtual Servers

All three components are required for every SLB configuration.

Nodes

A node is an individual service on a physical server that consists of an IP address and a port number.

Pools

A pool is a group of nodes that is mapped to a corresponding virtual server. Pools allow you to scale large networks that contain many nodes. Pools may be configured independently and associated with virtual servers in complex ways.

Each pool has its own load balancing method. When associated with a virtual server, the pool cannot be deleted from the SLB configuration.

Pools must be added before, and deleted after, the virtual servers that reference them. If a pool is not associated with a virtual server, it is not used for load balancing.

To create a pool, use this command:

```
create slb pool <poolname> {lb-method [round-robin
| ratio | priority | least-connections]}
```

To add nodes to a pool, use this command:

```
configure slb pool <poolname> add
<ipaddress>:<L4Port> {ratio <ratio> | priority
<priority>}
```

To delete nodes from a pool, use this command:

```
configure slb pool <poolname> delete
<ipaddress>:<L4Port>
```

Virtual Servers

Virtual servers are the backbone of the SLB configuration. They determine which groups of servers, or other network equipment, are targeted for server load balancing. Before you configure virtual servers, you need to know:

- The forwarding mode for your network design
- The name of the pool
- The virtual IP address

- The virtual port number

Once you know which virtual server options are useful in your network, you can:

- Define standard virtual servers
- Define wildcard virtual servers

Each virtual server maps to a single pool, which can be a group of content servers, firewalls, routers, or cache servers.

You can configure two different types of virtual servers:

- **Standard virtual servers**

  A standard virtual server represents a site, such as a Web site or an FTP site, and it provides load balancing for content servers. The virtual server IP address should be the same IP address that you register with the DNS (domain name system) for the site that the virtual server represents.

- **Wildcard virtual servers**

  A wildcard virtual server load balances transparent network devices such as firewalls, routers, or cache servers. Wildcard virtual servers use a special wildcard IP address (0.0.0.0), and you can use them only if Transparent mode is activated.

For cache server applications, use flow redirection.

A virtual server is identified by a virtual IP address. To create a virtual server, use this command:

```
create slb vip <vipname> pool <poolname> mode
[transparent | translation | port-translation]
<ipaddress>{-<upper_ipaddress>}: <L4Port> {unit
<number>}
```

Forwarding Modes

The 480T routing switch supports these SLB forwarding modes:

- Transparent
- Translational
- Port Translation
- GoGo

Table 18.1 summarizes the features supported by each forwarding mode.

Table 18.1: Forwarding Mode Feature Summary

	Transparent	Translational	Port Translation	GoGo
Performance	Hardware-based, server-to-client	Microprocessor-based, bi-directional	Microprocessor-based, bi-directional	Hardware-based, bi-directional
Load sharing algorithms	Round-robin, Ratio, Priority, Least Connections	Round-robin, Ratio, Priority, Least Connections	Round-robin, Ratio, Priority, Least Connections	Round-robin (hash)
Persistence	IPSA + Mask, IP list	IPSA + Mask, IP list	IPSA + Mask, IP list	IPSA
Health checking	L3, L4, L7, External	L3, L4, L7, External	L3, L4, L7, External	L1

Transparent Mode

As with any server load balancing application, the content must be duplicated on all physical servers.

Using transparent mode, the 480T routing switch does not modify the IP addresses before sending the traffic on to the selected server. To accomplish this, all servers must respond to the IP addresses associated with the virtual server. This virtual IP address (VIP) is the address used by the clients to connect to the virtual server. The servers must use this address as a loopback address and the address associated with the virtual server must be load balanced.

It is not possible to have a router between the SLB switch and the load balanced servers.

In transparent mode, servers can be directly attached or have a Layer 2 switch between the SLB switch and the server.Transparent mode is shown in Figure 18.1.

To configure transparent mode, use this command:

```
create slb vip <vipname> pool <poolname> mode
transparent <ipaddress>{-<upper_ipaddress>}:
<L4Port> {unit <number>}
```

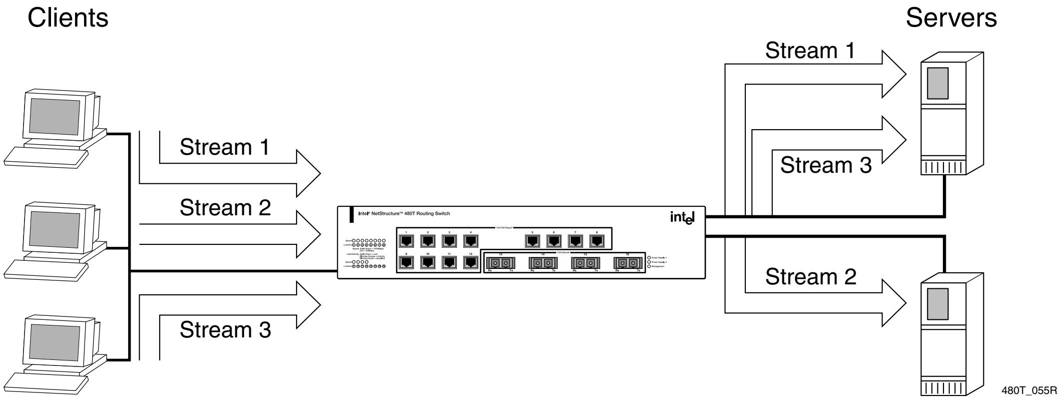

Figure 18.1: Transparent mode

In Figure 18.1, the 480T routing switch is configured to respond to requests for the VIP by forwarding them to the load balanced servers.

The servers are configured as follows:

- The interface for server 1 is 192.168.200.1
- The interface for server 2 is 192.168.200.2
- The loopback address on the servers is 192.168.201.1 (VIP)
- The service is configured to use the appropriate address and port, as specified in the switch configuration

The commands used to configure the switch as shown in Figure 18.1 are described below. These commands configure the VLANs and the switch IP addresses and subnets:

```
create vlan srvr

create vlan clnt

create vlan vips

configure srvr ipaddress 192.168.200.1 /24

configure clnt ipaddress 10.1.1.1 /24

configure vips ipaddress 192.168.201.1 /24

configure srvr add port 4-8

configure clnt add port 1-4

enable ipforwarding
```

Use these commands to create a round-robin pool called ***MyWeb***, and add nodes to the new pool:

```
create slb pool MyWeb lb-method round
```

```
configure slb pool MyWeb add 192.168.200.1:80
```

```
configure slb pool MyWeb add 192.168.200.2:80
```

Use this command to create a transparent mode VIP for the Web site and assign the ***MyWeb*** pool to it:

```
create slb vip WebVip pool MyWeb mode transparent
192.168.201.1:80
```

Use these commands to create a round-robin pool called ***MySSL***, and add nodes to the new pool.

```
create slb pool MySSL lb-method round-robin
```

```
configure slb pool MySSL add 192.168.200.1:443
```

```
configure slb pool MySSL add 192.168.200.2:443
```

This command creates a transparent mode ***VIP*** for the Web site and assigns the ***MySSL*** pool to it:

```
create slb vip SSLVip pool MySSL mode transparent
192.168.201.1:443
```

Use these commands to enable SLB, configure the server VLAN to act as the server side, and configure the client VLAN to act as the client side:

```
enable slb
```

```
configure vlan srvr slb-type server
```

```
configure vlan clnt slb-type client
```

Individual servers require that a loopback address be configured for each IP address to which the server will respond.

Translational Mode

In translational mode, requests coming in to the VIP are translated to the IP address of the server to be balanced. This mode does not require the configuration of a loopback address, since each server only uses its own IP address. As with any server load balancing application, the content must be duplicated on all physical servers.

To configure translational mode, use this command:

```
create slb vip <vipname> pool <poolname> mode
translation <ipaddress>{-<upper_ipaddress>}:
<L4Port> {unit <number>}
```

Figure 18.2 shows translational mode.

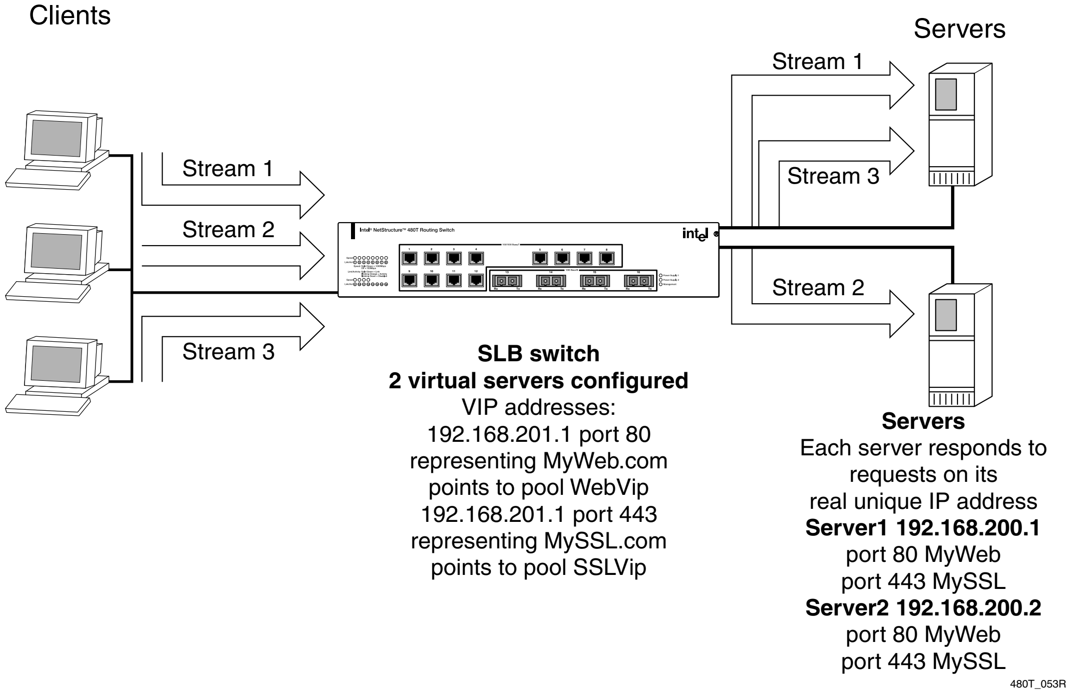

Figure 18.2: Translational mode

In Figure 18.2, the 480T routing switch is configured to respond to requests for the VIP by translating them and forwarding them to the load balanced servers. No additional server configuration is needed.

Use these commands to configure the VLANs and the switch IP addresses and subnets:

```
create vlan srvr

create vlan clnt

create vlan vips

configure srvr ipaddress 192.168.200.10 /24

configure clnt ipaddress 10.1.1.1 /24
configure vips ipaddress 192.168.201.1 /24
configure srvr add port 4-8
configure clnt add port 1-4
enable ipforwarding
```

These commands create a round-robin pool called ***MyWeb***, and add nodes to the new pool:

```
create slb pool MyWeb lb-method round
configure slb pool MyWeb add 192.168.200.1:80
configure slb pool MyWeb add 192.168.200.2:80
```

This command creates a translation mode VIP (virtual IP address) for the Web site and assign the ***MyWeb*** pool to it:

```
create slb vip WebVip pool MyWeb mode translation
192.168.201.1:80
```

Use these commands to create a round-robin pool called ***MySSL***, and add nodes to the new pool:

```
create slb pool MySSL lb-method round
configure slb pool MySSL add 192.168.200.1:443
configure slb pool MySSL add 192.168.200.2:443
```

To create a translation mode VIP for the Web site and assign the ***MySSL*** pool to it, use this command:

```
create slb vip SSLVip pool MySSL mode translation
192.168.201.1:443
```

Use these commands to enable SLB, configure the server VLAN to act as the server side, and configure the client VLAN to act as the client side:

```
enable slb
configure vlan srvr slb-type server
configure vlan clnt slb-type client
```

Port Translation Mode

Port translation is essentially the same thing as translational mode, except that the Layer 4 port on the virtual server can be different from the Layer 4 port on the nodes being load balanced. The 480T routing switch automatically changes the IP address and port address on incoming packets to that of the load balanced servers. As with any server load balancing application, the content must be duplicated on all physical servers.

Configure port translation mode using this command:

```
create slb vip <vipname> pool <poolname> mode
port-translation <ipaddress>-
{<upper_ipaddress>}:<L4Port> {unit <number>}
```

GoGo Mode

As with any server load balancing application, the content must be duplicated on all physical servers.

GoGo mode is considered a fast (line rate) method of server load balancing. GoGo mode forwards traffic without manipulating packet content. Session persistence is maintained using IP source address persistence information.

Traffic is optimally balanced across groups of two, four, or eight directly attached servers. Because servers are always directly attached, there is no need to configure nodes, pools, or VIPs.

To use GoGo mode, all servers are configured with the same MAC and IP addresses. In GoGo mode, the load balancing method is fixed, and is based on a hashing of the client IP address. All GoGo mode traffic exhibits persistence based on source IP information. That is, a given source address is mapped to one and only one physical server.

Figure 18.3 shows GoGo mode.

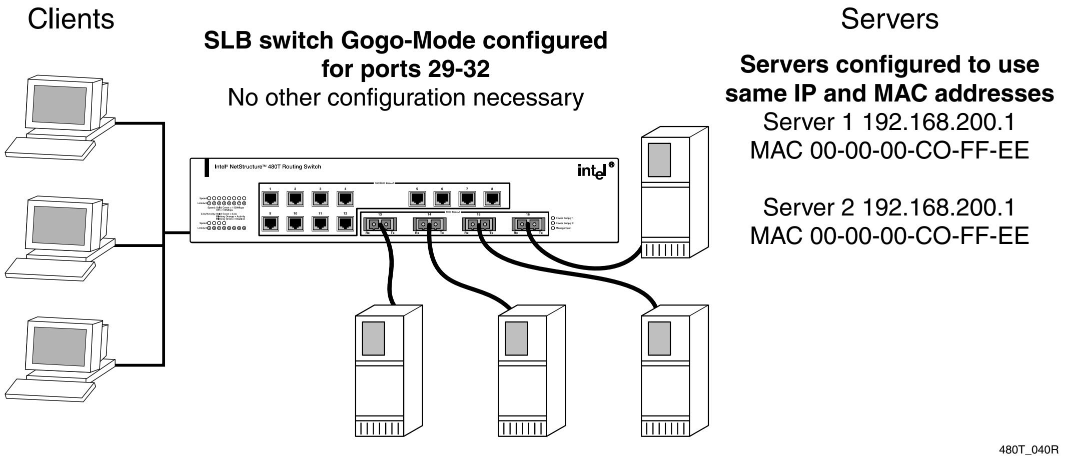
Clients
SLB switch Gogo-Mode configured for ports 29-32
No other configuration necessary
Servers
Servers configured to use same IP and MAC addresses
Server 1 192.168.200.1
MAC 00-00-00-CO-FF-EE
Server 2 192.168.200.1
MAC 00-00-00-CO-FF-EE
Intel® NetStructure™ 480T Routing Switch
intel
480T_040R

Figure 18.3: GoGo mode

In Figure 18.3, the 480T routing switch is configured to balance all traffic sent to the VIP based on the client IP address.

All servers have the same:

- MAC address
- IP address
- Content

The commands used to configure the switch, as indicated in the example are:

```
create vlan server

create vlan client

configure srvr ipaddress 10.1.1.1 /24

configure clnt ipaddress 1.1.1.1 /24

configure srvr add port 4-8

configure clnt add port 1-4

enable slb gogo 4 grouping 4-8

enable ipforwarding
```

Separating clients and servers into separate VLANs is not a requirement in GoGo mode.

VIP Network Advertisement

There are three methods for controlling network connectivity to the VIPs. Depending on the subnet to which the VIP belongs, the 480T routing switch will adjust its behavior automatically.

- **Proxy ARP -** If the VIP is a member of an existing subnet to which the switch is directly attached, the switch responds to ARP requests on behalf of the VIP. This allows you to implement server load balancing on a Layer 2 network. The VLAN containing the servers is a different subnet than the client VLAN's subnet. The VIP appears as a member of the client subnet.
- **Host-Route -** If the VIP created is not a member of an existing subnet that the switch is directly attached to, a host-route entry is added to the routing table for the switch. All clients will need to have a routed path to the VIP that points to the switch's IP address on the client VLAN.

- **Subnet-Route -** If your network configuration requires that the VIPs be propagated through a routing protocol by the switch, you need to create a loopback VLAN with the VIP(s) being valid members of the loopback VLAN’s subnet. When a routing protocol is enabled, the subnet containing the VIPs is propagated through the network.

Balancing Methods

A load balancing method defines, in part, the logic that the 480T routing switch uses to determine which node should receive a connection hosted by a particular virtual server. Individual load balancing methods take into account one or more dynamic factors, such as current connection count.

Because each application of SLB is unique, node performance depends on a number of different factors. We recommend that you experiment with different load balancing methods, and choose the one that offers the best performance in your particular environment.

The 480T routing switch supports these load balancing methods:

- Round-robin
- Ratio
- Least connections
- Priority

Round-Robin

The default load balancing method is round-robin, and it simply passes each new connection request to the next server in line, eventually distributing connections evenly across the array of devices being load balanced. Round-robin works well in most configurations, especially if the equipment that you are load balancing is roughly equal in processing speed and memory.

To configure round-robin, use this command:

```
configure slb pool <poolname> lb-method round-robin
```

Ratio

If you are working with servers that differ significantly in processing speed and memory, you may want to switch to the ratio load balancing method. In ratio, the 480T routing switch distributes connections among devices according to ratio weights that you set, where the number of connections that each device receives over time is proportionate to the ratio weight.

For example, if your array contained one new, high-speed server and two older servers, you could set the ratio so that the high-speed server receives twice as many connections as either of the two older servers.

To configure ratio, use this command:

```
configure slb pool <poolname> lb-method ratio
```

Ratio Weight

The ratio weight is the proportion of total connections that the node address should receive. The default ratio weight for a given node address is 1. If all node addresses use this default weight, the connections are distributed equally among the nodes. A ratio weight of 2 would result in twice as much traffic as a ratio weight of 1.

To configure a ratio weight, use this command:

```
configure slb pool <poolname> add
<ipaddress>:<L4Port> ratio {<ratio>}
```

Least Connections

The least connections method is considered relatively simple in that the switch passes a new connection to the node having the least number of active sessions. The number of active sessions includes only those sessions occurring within the same VIP (virtual IP address). Least connections works best in environments where the servers or other equipment you are load balancing have similar capabilities.

To configure least connections, use this command:

```
configure slb pool <poolname> lb-method least-
connections
```

Priority

Priority mode is a variant of round-robin designed to provide redundant standby nodes within a pool. When you add a node to a pool, you can assign a priority level. Priority numbers range from 1 to 65,535, with the highest number indicating the highest priority.

The 480T routing switch will distribute traffic in round-robin fashion among the pool's active nodes with the highest priority. If all nodes at that priority level go down or hit a session limit maximum, all new sessions are directed to the nodes at the next lowest priority level.

The switch continually monitors the status of the down nodes. As each node comes back up, the switch distributes traffic according to the priorities.

For example, with a pool that has six nodes divided evenly into two priority levels (2 and 1) all sessions are evenly distributed using round-robin to the nodes at priority level 2.

If one of the priority level 2 nodes are down, all of the traffic is assigned to the remaining level 2 nodes. If all of the priority level 2 nodes are down, then all sessions are directed to the priority level 1 nodes. If one of the level 2 nodes comes back up, all new sessions are assigned to it.

Basic SLB Commands

Table 18.2 describes basic SLB commands. Press the Tab key in the command line interface for further command options.

Table 18.2: Basic SLB Commands

Command	Description
clear slb connections [VIP <vipname> | ipaddress <ipaddress>:{<L4Port>}]	Clears the active connections.
configure slb pool <poolname> add <ipaddress>:<L4Port> {ratio <ratio> | priority <priority>}	Adds a physical server (node) to a server pool. When a new node is added, **`ping-check`** is automatically enabled.

Table 18.2: Basic SLB Commands

Command	Description
configure slb pool <poolname> delete <ipaddress>:<L4Port>	Deletes a physical server from a server pool.
configure slb pool <poolname> lb-method [round-robin | ratio | priority | least-connections]	Configures the SLB load-balancing method.
configure slb l4-port <L4Port> [treaper_timeout <seconds> | udp-idle-timeout <seconds>]	Configures the inactive period for TCP or UDP before the connection is aged out.
configure vlan <name> slb-type [server | client | both | none]	Marks a VLAN as either a server VLAN or a client VLAN. If the server also originates connections to other servers, set **`slb-type`** to both.
create slb pool <poolname> {slb-method [round-robin | ratio | priority | least-connections]}	Creates a server pool and optionally assigns a load-balancing method to the pool. The default load-balance method is round-robin. A pool represents a group of physical servers that is used to load-balance one or more VIPs.
create slb vip <vipname> pool <poolname> mode [transparent | translation | port-translation] <ipaddress> {-<upper_ipaddress>}:<L4Port> {unit <number>}	Creates one or more new virtual IP addresses (VIPs) and attaches the VIP to a pool of physical servers. The server pool needs to be created before the VIP is created. If **`port`** is not specified, all requests to the VIP are forwarded to the server. If **`port`** is specified, only the specified TCP/UDP ports are allowed to reach the server. All other packets are dropped.
delete slb pool [<poolname> | all]	Deletes a server pool.
delete slb vip [<vipname> | all]	Deletes one or all VIPs.

Table 18.2: Basic SLB Commands

Command	Description
disable slb	Disables SLB processing. Disabling SLB: • Closes all connections. • Withdraws VIP routes or routes that do not respond with proxy ARP responses of VIP addresses. • Disconnects the switch from redundant SLB switches.
disable slb gogo-mode <port number> {all | ping-check | tcp-port-check <port> | service-check <port>}	Disables gogo-mode processing.
disable slb node [<ipaddress>:<L4Port> | all] {close-connections-now | [ping-check | tcp-port-check]}	Disables one or more nodes from receiving new connection establishments. If `close-connections-now` is specified, all current open connections are closed immediately.
disable slb l4-port [<L4Port> | all]	Disables one or all L4 ports for SLB.
disable slb vip ipaddress <ipaddress>:<L4Port> {close-connections-now}	Disables a single VIP port.
disable slb vip <vipname> {close-connections-now}	Disables a VIP group. When disabled, no new connections to the real servers are allowed. If `close-connections-now` is specified, all existing connections are immediately closed. Otherwise, the existing connections are closed naturally, and are subject to connection reaping if idle for longer than the treaper-timeout configured on the SLB port.
disable slb vip all {close-connections-now | client-persistence | service-check | sticky-persistence | svcdown-reset}	Disables all VIP groups. If `close-connections-now` is specified, all existing connections are immediately closed.

Table 18.2: Basic SLB Commands

Command	Description
enable slb	Enables SLB processing on the switch, and activates these functions for transparent, translational, and port translation modes: • Exporting of VIP routes or proxy ARP for VIP addresses. • Processing of VIP lookup and connection setup. • Establishing communication with redundant SLB switches. The default setting is disabled.
enable slb gogo-mode <port number> grouping <portlist>	Enables gogo-mode processing for a group of ports. There are no additional configuration commands for gogo mode.
enable slb node [<ipaddress>:<L4Port> | all] ping-check | tcp-port-check}	Enables one or more nodes to receive data traffic. A node represents a physical server.
enable slb l4-port <L4Port>	Enables an L4 port for SLB.
enable slb vip ipaddress <ipaddress>:<L4Port>	Enables a single VIP port.
enable slb vip <vipname>	Enables a VIP group.
show slb	Displays the current SLB global configuration information, including: • Global enable/disable mode • Global modes • Default settings for health checker
show slb node {<ipaddress>:<L4Port>}	Displays node-specific configuration information and status.
show slb pool {poolname}	Displays the current SLB pool configuration and statistics.

Table 18.2: Basic SLB Commands

Command	Description
show slb pool <poolname>	Displays the configuration for the specified SLB pool.
show slb l4-port {<L4Port>}	Displays the SLB configuration for one or all L4 ports.
show slb vip {detail}	Displays the current VIP configuration and statistics.
show slb vip <vipname> {detail}	Displays the configuration for the specified VIP.
unconfigure slb all	Resets SLB global defaults and clears the SLB configuration.

Advanced SLB Application Example

This example builds upon the introductory SLB example. The advanced concepts included in this example are:

- Multiple pools
- Multiple VIPs
- Multiple balancing algorithms
- Multiple types of health checking

Figure 18.4 shows an example of an advanced SLB application.

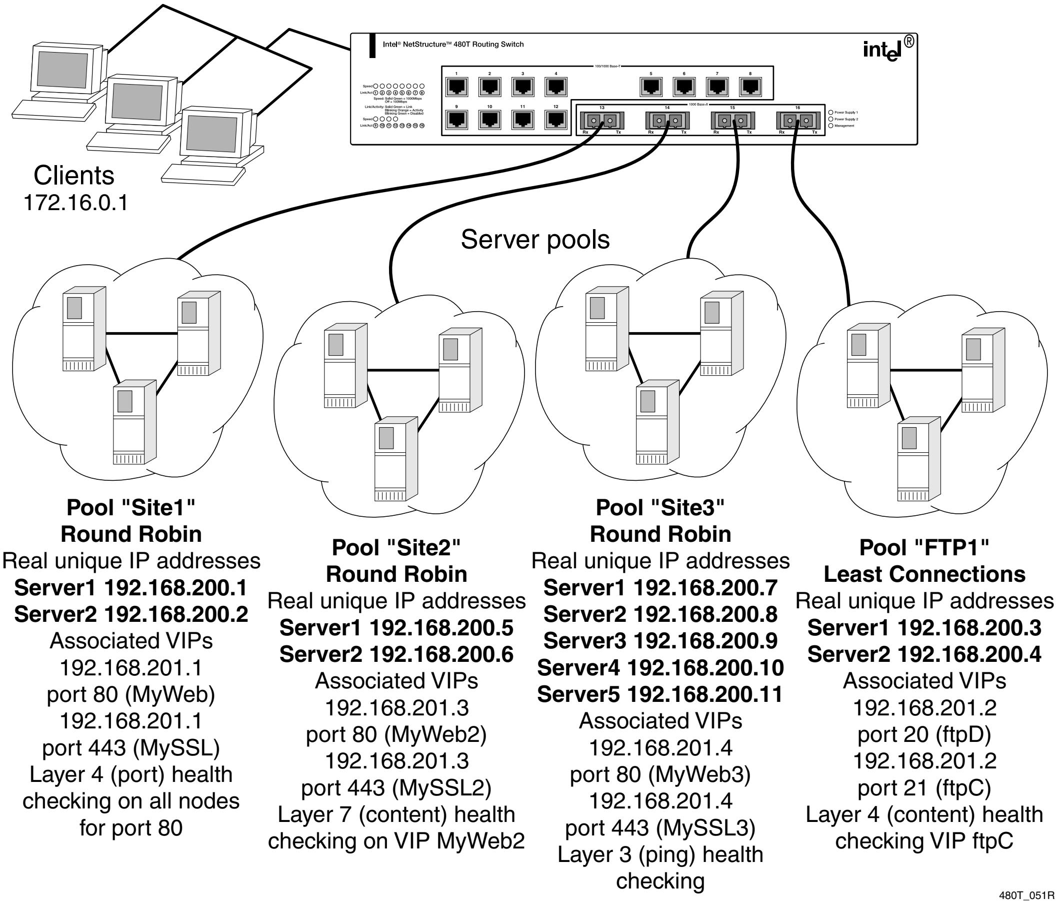

Figure 18.4: Advanced SLB configuration

The commands used to configure are described below.

Use these commands to create the VLAN from which outside connections will come:

```
create vlan outside
```

```
configure vlan outside ipaddress 172.16.0.1 /16
```

```
configure vlan outside add ports 1-8
```

To create the virtual IP VLAN, use these commands:

```
create vlan sites
```

```
configure vlan sites ipaddress 192.168.201.254 /24
```

All VIPs is configured to use this subnet. There are no ports associated with this VLAN.

You can use these commands to create the VLAN ***servers*** and enable IP forwarding:

```
create vlan servers

configure vlan servers ipaddress 192.168.200.254 /
24

configure vlan servers add ports 9-16

enable ipforwarding
```

The next example shows a series of commands used to create a Web site. The site is defined as having 2 servers: 192.168.200.1 and 192.168.200.2, each with 2 services (HTTP and SSL).

Two VIPs (virtual IP addresses) are then created to point at the appropriate pools. As a default, round-robin is used to load balance the services.

Only one IP address is used for both VIPs; the difference is the port number. Finally, port checking is enabled to ensure fault tolerance on each of the servers.

Use these commands:

```
create slb pool site1web

configure slb site1 add 192.168.200.1:80

configure slb site1 add 192.168.200.2:80

create slb pool site1ssl

configure slb site1 add 192.168.200.1:443

configure slb site1 add 192.168.200.2:443

create slb vip myweb pool site1web mode
transparent 192.168.201.1:80

create slb vip myssl pool site1ssl mode
transparent 192.168.201.1:443

enable slb node 192.168.200.1:80 tcp-port-check

enable slb node 192.168.200.2:80 tcp-port-check

enable slb node 192.168.200.1:443 tcp-port-check
enable slb node 192.168.200.2:443 tcp-port-check
```

The next series of commands creates a second Web site. This second site is similar to the first example; the difference is that content checking is enabled on this site. For this type of health checking, the server downloads a specified page (/testpage.htm) and looks for a specific string in the content (`“test successful”`). If it finds the string, the server is considered online.

Use these commands:

```
create slb pool site2web

configure slb site2web add 192.168.200.5:80

configure slb site2web add 192.168.200.6:80

create slb pool site2ssl

configure slb site2ssl add 192.168.200.5:443

configure slb site2ssl add 192.168.200.6:443

create slb vip myweb2 pool site2web mode
transparent 192.168.201.3:80

create slb vip myssl2 pool site2ssl mode
transparent 192.168.201.3:443

enable slb vip myweb2 service-check

configure slb vip myweb2 service-check http url
“/testpage.htm” match-string “test successful”
```

The following series of commands creates a third Web site. This example creates one pool with a wildcard port specified. This means that the pool allows any port that is sent to it by the VIP. All five servers respond to requests on both port 80 and port 443.

Note: Port 0 the wildcard port.

Use these commands:

```
create slb pool site3web

configure slb site3web add 192.168.200.7:0

configure slb site3web add 192.168.200.8:0

configure slb site3web add 192.168.200.9:0

configure slb site3web add 192.168.200.10:0

configure slb site3web add 192.168.200.11:0

create slb vip myweb3 pool site3web mode
transparent 192.168.201.4:80
```

```
create slb vip myssl3 pool site3web mode
transparent 192.168.201.4:443
```

The next example demonstrates the series of commands you would use to create an FTP site.

The site is defined as having two servers: 192.168.200.3 and 192.168.200.4. Only FTP is being serviced by the servers. The two different VIPs and port numbers refer to the control and data channels used by the FTP service. Two VIPs are then created to point at the appropriate pools.

As with the first site, the default load balancing method (round-robin) is used. Layer 7 health checking is used on the ftpC VIP. By using health checking, the switch logs in to the site as user ***test*** with the password ***testpass***.

If the login is successful, the server is labeled up and is allowed to participate in load balancing. The account and password must be set up on all FTP servers. Use these commands:

```
create slb pool ftp1c

configure slb ftp1c add 192.168.200.3:21

configure slb ftp1c add 192.168.200.4:21

create slb pool ftp1d

configure slb ftp1d add 192.168.200.3:20

configure slb ftp1d add 192.168.200.4:20

create slb vip ftpc pool ftp1c mode transparent
192.168.201.2:21

create slb vip ftpd pool ftp1d mode transparent
192.168.201.2:20

enable slb vip ftpc service-check

configure slb vip ftpc service-check ftp user test
password testpass
```

Finally, enable SLB and configure the VLANs as either client or server, using these commands.

```
enable slb

configure vlan outside slb-type client

configure vlan servers slb-type server
```

Health Checking

The 480T routing switch supports both internal and external health checking.

Health check definitions

For reference, the following health checks are available on all Server Load Balancing, Web Cache Redirection and Policy-based Routing functions. SLB functions test individual servers. Web Cache Redirection and Policy-based routing functions test the next-hops in accordance with the flow-redirection rules.

Layer 3 Ping Check

The default health checking is a simple ping check, where the switch sends an ICMP ping packet to the configured server or next-hop. If 3 replies are lost, the server or next-hop is set to down and flows are not redirected to it. The ping check is the only health checking that will work with a wildcard as the IP Port.

Layer 4 Port Check

The switch will attempt to establish a TCP connection to the server or next-hop.

When using Web Cache Redirection or Policy Based Routing, the Layer 4 port must be defined in the flow and opened on the next-hop in order for the health check to succeed.

Layer 7 HTTP Check

The HTTP health check will download a specific page from the server or next-hop configured for the flow. The switch will then search the page for a specific text string in the first 1000 bytes. If the text string is found, the check passes. As an alternative you can configure the check to accept any data from the downloaded page.

Layer 7 FTP Check

The FTP health check establishes an FTP connection between the switch and the server or next-hop. The switch will attempt to log in using the name and password supplied during the configuration. The check will succeed when the switch successfully logs into the next-hop.

Layer 7 NNTP Check

The NNTP health check connects to the server or next-hop, establishes a connection, and attaches to a user defined newsgroup.

Layer 7 POP3, SMTP, and Telnet Check

These health checks attach to the server or next-hop using the specified protocol and log in. After successful login the next-hop is marked as up.

Internal Health Checking

Three types of internal health checks are available:

- Ping-check
- Port-check
- Service-check

If any of the health checks enabled on a given node do not pass within the timeout specified, the node is considered down. When a node is down, no new connection is established to that node until the node passes all configured health checks. If a health check fails and if the svcdown-reset parameter is enabled on an associated VIP (virtual IP address), existing connections for the VIP on this node is closed by sending TCP Reset to the client and node.

In the command-line interface, the commands **`show pool`** and **`show vip`** display individual node resources as up or down. New connections are only allowed if the VIP and node in question are both enabled and up. A node is assumed to be up unless health check is enabled and fails, in which case the node is marked down.

A resource is also marked down if it was disabled and the number of existing connections drops to zero. If a node is marked down for this reason, ping-checks and port-checks on this node are automatically stopped to conserve system resources, but they resume if the node is enabled by the user.

The 480T routing switch also supports external health checking. External health checking uses an external service configured by the user to perform health checks and uses SNMP (Simple Network Management Protocol) as a mechanism to notify the switch of a server failure.

Ping-Check

Ping-check is Layer 3-based pinging of the physical node. The default ping frequency is one ping generated to the node every 10 seconds. If the node does not respond to any ping within a timeout period of 30 seconds (3 ping intervals), the node is considered down.

To enable ping-check, use this command:

```
enable slb node <ipaddress> ping-check
```

To disable ping-check, use this command:

```
disable slb node <ipaddress> ping-check
```

TCP-Port-Check

TCP-port-check is Layer 4-based TCP port open/close testing of the physical node. The default frequency is 30 seconds and the default timeout is 90 seconds.

Port-checking is useful when a node passes ping-checks, but a required TCP service (for example, HTTP) has gone down. If the HTTP service running on TCP port 80 crashed, that would cause a Layer 4 port-check on port 80 to fail, because no TCP socket could be opened to that port. If this continues for the duration of the specified port-check timeout, the IP/port combination is considered down.

To enable tcp-port-check, use this command:

```
enable slb node <ipaddress>:<L4Port> tcp-port-check
```

To disable tcp-port-check, use this command:

```
disable slb node <ipaddress>:{<L4Port> | all} tcp-
port-check
```

Service-Check

Service-check is Layer 7-based and application-dependent. It is defined on a VIP and is performed on each node in the pool with which this VIP is associated. The default frequency is 60 seconds and the default timeout is 180 seconds. Each service check has associated parameters that you can set. These parameters are described in Table 18.3.

Table 18.3: Service-Check Parameters

Service	Attribute	Global Default Value
HTTP	URL Match-string	/ Any-content
FTP	Userid Password	anonymous anonymous
Telnet	Userid Password	anonymous anonymous
SMTP	DNS-domain	Same as the switch DNS domain. If no DNS domain is configured for the switch, the value is *mydomain.com*.
NNTP	Newsgroup	ebusiness
POP3	Userid Password	anonymous anonymous

If the service-check parameters are not specified on an individual node or VIP, the global default values for these parameters are used. The global defaults are configurable, so you can change them to your most often used parameters.

In the case of HTTP service-checking, you can specify the URL of the Web page to be retrieved, such as /index.html. You can also specify a match-string, such as Welcome, that is expected to be in the retrieved Web page.

If the match-string is found in the first 1,000 bytes of the retrieved Web page, the service-check passes on the particular node. A match-string specifying the keyword any-content will match any retrieved text. However, to distinguish valid data in the retrieved text from error text, we recommend that you specify an actual string to match.

For FTP, Telnet, and POP3, service-check attempts to log on and off the application on the server using the specified userID and password.

For SMTP, service-check identifies the identity of the switch by providing the specified DNS domain. The SMTP server might not even use the specified DNS domain for authentication, only identification.

For NNTP, service-check queries the newsgroup specified.

Because service-checking is configured on a VIP basis, multiple VIPs can use the same nodes, and you can run multiple service-checks against a particular node IP address and port number, it is possible for some of these service-checks to fail, while others pass. Therefore, when determining if a given node can accept a new connection for a VIP, the node must pass the service-check configured for that VIP.

When showing detailed VIP information, the status for individual nodes is shown with respect to that VIP.

To enable service-check, use this command:

```
enable slb vip [<vipname> | all] service-check
```

To disable service-check, use this command:

```
disable slb vip [<vipname> | all] service-check
```

GoGo Mode Health Checking

The 480T routing switch supports health checking on servers participating in SLB GoGo mode. You can configure multiple health checks (ping-check, tcp-port-checks and service-checks) simultaneously on a given GoGo mode grouping. A physical port in a GoGo mode grouping is considered available for GoGo traffic only if all configured health checks pass.

Use these commands to enable GoGo mode health checking:

```
enable slb gogo-mode master ping-check {ipaddress}
```

```
enable slb gogo-mode master tcp-port-check [port |
all]
```

```
enable slb gogo-mode master service-check [http | ftp
| telnet | smtp | nntp | pop3 | all | tcpport]
```

Use these commands to disable GoGo mode health checking:

```
disable slb gogo-mode master ping-check
```

```
disable slb gogo-mode master tcp-port-check [port |
all]
```

```
disable slb gogo-mode master service-check [http |
ftp | telnet | smtp | nntp | pop3 | all | tcpport]
```

```
unconfigure slb gogo-mode master health-check
```

This command disables and deletes all ping-check, tcp-port-check, and service-check configurations for this GoGo mode grouping. The GoGo mode grouping itself is not affected.

```
unconfigure slb gogo-mode master service-check [http
| ftp | telnet | smtp | nntp | pop3 | all | tcpport]
```

This command disables and deletes the service-check configuration. If the associated TCP port has not been used for any tcp-port-check configuration, the TCP port is deleted as well.

Use these commands to configure GoGo mode health checking:

```
configure slb gogo-mode master ping-check frequency
seconds timeout seconds
```

```
configure slb gogo-mode master health-check ipaddress
```

```
configure slb gogo-mode master tcp-port-check [add |
delete] port
```

```
configure slb gogo-mode master tcp-port-check timer
port frequency seconds timeout seconds
```

```
configure slb gogo-mode master service-check http
{l4-port port} {url url match-string [match_string |
any-content]}
```

```
configure slb gogo-mode master service-check ftp {l4-
port port} {userid userid | password {encrypted}
password}
```

```
configure slb gogo-mode master service-check telnet
{l4-port port} {userid userid | password {encrypted}
password}
```

```
configure slb gogo-mode master service-check smtp
{l4-port port} {dns_domain}
```

```
configure slb gogo-mode master service-check nntp
{l4-port port} {newsgroup}
```

```
configure slb gogo-mode master service-check pop3
{l4-port port} {userid userid password {encrypted}
password}
```

```
configure slb gogo-mode master service-check timer
[http | ftp | telnet | smtp | nntp | pop3 | tcpport]
frequency seconds timeout seconds
```

Use these command to view your GoGo mode health checking configuration:

```
show slb gogo-mode {master} {configuration}
```

SLB Global Connection Timeout

For SLB transparent and translational modes you can configure the global connection timeout period. This helps to avoid cases where connections are closed because the TCP FIN and ACK timeout is too short.

To configure the global connection timeout period (between 1 and 180 seconds) use this command:

```
configure slb global connection-timeout <seconds>
```

The default value is 1 second. In addition, the timeout should be set as low as possible to avoid stale connections staying in the table.

External Health Checking

For server health checking, that goes beyond the abilities of internal health checking, the 480T routing switch also supports external health checking. The external health checking device sends the results of its check to the switch by way of SNMP MIB attributes. For information on the specific MIB definitions for external health checking, contact Intel Customer Support (see Appendix D, "Intel Customer Support" on page 461).

Health Checks for Web Cache Redirection and Policy Based Routing

Health checking works on the ports configured by their associated flow. For example, if you configure a flow to redirect on port 80 (HTTP), but FTP is configured as the service check, the switch will try to open an FTP session on port 80. The health check will fail if the protocol will not work on the configured flow.

Several additional health checks are supported for the flows that are defined under web cache redirection and policy based routing. The operation and definition of these health checks are identical to those used for server load balancing.

- **Ping Check:** The ping check is the only health checking that will work with a wildcard as the Layer 4 IP port. To configure a ping check for a defined flow, use this command:

```
configure <flow> service-check ping
```

- **Layer 4 Port Check:** The port has to be defined and open on the next hop in order for the health check to succeed. To configure a Layer 4 Port health check for a defined flow, use this command:

```
configure <flow> service-check L4-port
```

- **HTTP Check:** To configure an HTTP health check for a defined flow, use this command:

```
configure <flow> service-check http url "<url>"
match-string "<string>"
```

In this example the switch will connect to the cache and download the page test.htm in the root WWW directory and search the page for the word pass in the first 1000 bytes. The quotation marks are necessary for the switch to recognize the Web page and the string.

- **FTP Check:** To configure an FTP health check for a defined flow, use this command:

```
configure <flow> service-check ftp user <user>
<password>
```

- **NNTP Check:** To configure an NNTP health check for a defined flow, use the command:

```
configure <flow> service-check nntp <newsgroup>
```

- **POP3, SMTP and Telnet Checks:** To configure a POP3, SMTP or Telnet health check for a defined flow, use the command:

```
configure <flow> service-check <pop3|smtp|telnet>
user <user> <password>
```

Configuring health check timeouts and frequencies is similar to the server load balancing command:

```
conf slb global service-check frequency <seconds>
timeout <seconds>
```

Layer 4 Flows

Policy-based routing and Web cache redirection support an **any** option for the Layer 4 protocol type which allows the redirection of TCP, UDP and other traffic types with the exception of ICMP traffic. To configure this capability, use the **any** option in the syntax for flow re-direction.

```
create flow-redirect <flow_rule_name> [tcp | udp |
any] destination [<ip_address>/<mask> | any] ip-port
[<L4_port> | any] source [<ip_address>/<mask> | any]
```

Policy-Based Routing with Route Load-Sharing

Policy-based routing is used to alter the normally calculated next-hop route which is based on the route table. This same alteration can also load-share across multiple routers. It implies a set of rules or policies that take precedence over information in the route table. These policies can perform a flow-redirection to different next-hop addresses based on:

- IP source address and mask
- IP destination address and mask

Layer 4 Destination Port

In the event that the next-hop address (or addresses) becomes unavailable, the switch will route the traffic normally. Several rules can be defined; the precedence of rules is determined by best match of the rule to the packet. If no rule is satisfied, no redirection occurs.

There are two types of commands to setup policy-based routing, one to configure the redirection rules and one to configure the next-hop IP addresses:

```
create flow-redirect <flow_rule_name> [tcp | udp |
any] destination [<ip_address>/<mask> | any] ip-port
[<L4_port> | any] source [<ip_address>/<mask> | any]
```

```
configure flow-redirect <flow_rule_name> [add |
delete] next-hop <ip_address>
```

If multiple next-hop addresses are defined, traffic satisfying the rule is load-shared across the next-hop addresses based on destination IP address. If next-hop addresses fail (do not respond to ICMP pings), the switch will resume normal routing. Using policy-based routing has no impact on switch performance.

To show configuration and status of flow redirection rules, use this command:

```
show flow-redirect [<flow_rule_name | <cr>]
```

Maintenance Mode

You can easily put a node or VIP into maintenance mode by disabling the node or VIP. In maintenance mode, existing connections remain active, but no new connections are permitted. The existing connections are either closed by the client and server, or are aged out if idle for more than 600 seconds.

Persistence

Using persistence, you can ensure that traffic flows do not span multiple servers. The 480T routing switch supports two types of persistence:

- Client persistence
- Sticky persistence

Client Persistence

Client persistence for a virtual server provides a persist mask feature. You can define a range of IP addresses that can be matched to a persistent connection. Any client whose source IP address falls within the range is considered a match for the given persistence entry.

To configure client persistence, use this command:

```
enable slb vip [<vipname> | all] client-
persistence {timeout <seconds>} {mask <mask>}
```

SLB Proxy Client Persistence

Use SLB proxy client persistence when you need client persistence and you use multiple NAT address ranges to translate the internal client IP addresses. Use these three commands:

```
enable slb proxy-client-persistent
```

```
disable slb proxy-client-persistent
```

```
configure slb proxy-client-persistent [add |
delete] <ipaddress / mask>
```

Sticky Persistence

Sticky persistence provides a special type of persistence that is especially useful for cache servers. Similar to client persistence, sticky persistence keeps track of incoming clients' source and destination IP addresses.

When a client is looking to make a repeat connection to a particular destination IP address, the 480T routing switch directs the client to the same cache server or other transparent node that it used previously.

Allowing clients to repeatedly use the same cache server can help you reduce the amount of content that might otherwise be duplicated on two or more cache servers in your network.

To prevent sticky entries from clumping on one server, use a static load balancing mode, such as round-robin.

You can only activate sticky persistence on wildcard virtual servers.

Sticky persistence provides the most benefit when you load balance caching proxy servers. A caching proxy server intercepts Web requests and returns a cached Web page if it is available. To improve the efficiency of the cache on these proxies, it is necessary to send similar requests to the same proxy server repeatedly.

You can use sticky persistence to cache a given Web page on one proxy server instead of on every proxy server in an array. This saves the other proxies from duplicating the Web page in their cache, wasting memory.

To configure sticky persistence, use this command:

```
enable slb vip [<vipname> | all] sticky-
persistence {timeout <seconds>}
```

Server Load Balancing with ESRP

Using ESRP (Enterprise Standby Router Protocol), the SLB service is made redundant, along with the Layer 2 and Layer 3 services of the 480T routing switch . This configuration allows single- or dual-attached servers to support redundant gateway services and very fast recovery from a fault.

When ESRP is enabled, all servers can be online at the same time (as opposed to only the ones connected to the active switch in High

Availability mode or having to introduce another interconnecting switch), and recovery from a switch failure occurs in less than 8 seconds.

Figure 18.5 shows SLB enabled using ESRP and dual-attached servers.

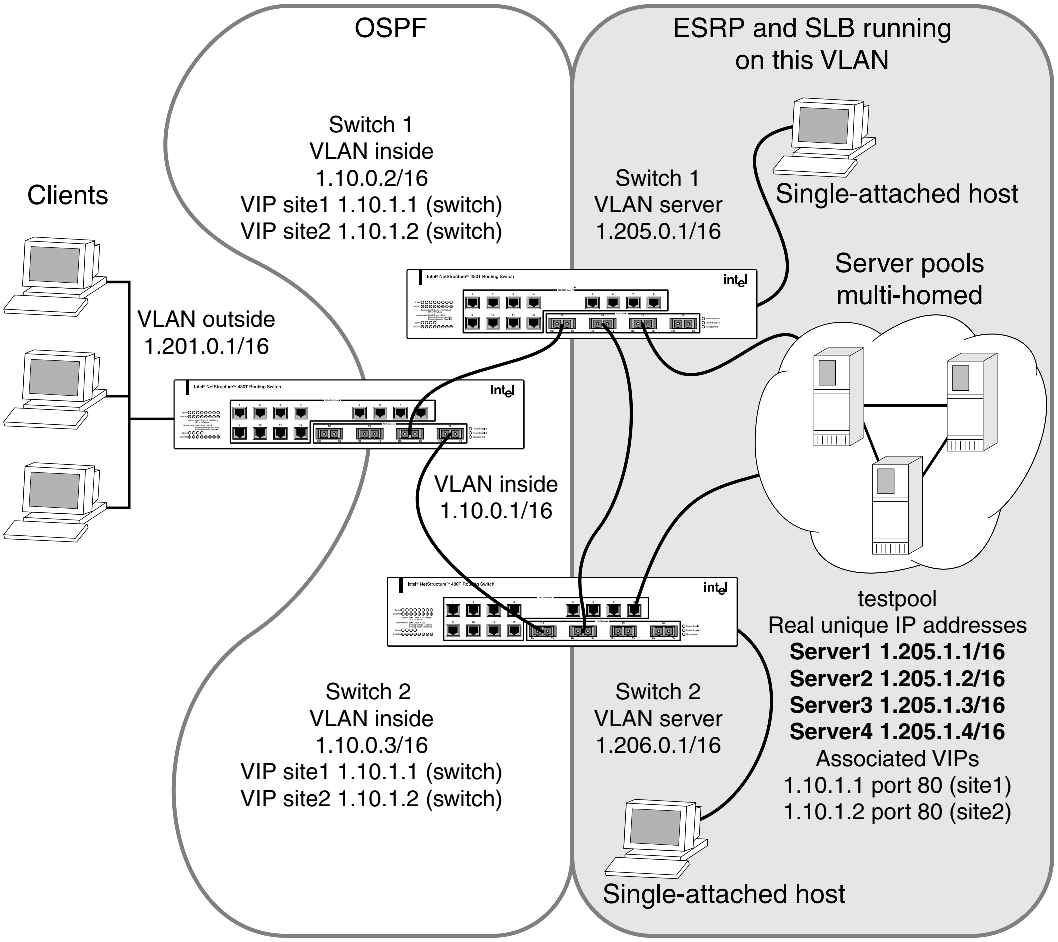

480T physical configuration
VLAN outside
Dual-attached servers connected to ports 1-4
Interconnect (also configured as host) on port 15
VLAN inside
port 16 connects to gateway switch

480T physical configuration
VLAN server
Dual-attached servers connected to ports 1-4
Interconnect (also configured as host) on port 32
VLAN inside
port 31 connects to gateway switch

Figure 18.5: SLB using ESRP and dual-attached servers

Configuring the Switches for SLB and ESRP

The SLB and ESRP configurations are identical on both switches, in relation to the ports being used.

The procedure used to configure the Switch 1 and Switch 2 in Figure 18.5 is described below.

1. Create the VLANs, using these commands:

   ```
   create vlan inside
   create vlan server
   ```

2. Connect the gateway to the VLAN ***inside***, using these commands:

   ```
   configure inside ipaddress 1.10.0.2 /16
   configure inside add port 10
   ```

3. Configure the servers to connect to the VLAN ***server*** on ports 1 through 4, and configure port 8 to connect to the other ESRP switch, using these commands:

   ```
   configure server ipaddress 1.205.0.1 /16
   configure server add port 1-4, 8
   ```

4. Enable IP forwarding, create a server pool called ***testpool,*** and add four servers to it using TCP port 80, using these commands:

   ```
   enable ipforwarding
   create slb pool testpool
   configure slb pool testpool add 1.205.1.1:80
   configure slb pool testpool add 1.205.1.2:80
   configure slb pool testpool add 1.205.1.3:80
   configure slb pool testpool add 1.205.1.4:80
   ```

5. Use these commands to create SLB VIP addresses for the two Web sites (***site1*** and ***site2***) and associate them with server pool ***testpool***:

   ```
   create slb vip site1 pool testpool mode
   transparent 1.10.1.1:80
   create slb vip site2 pool testpool mode
   transparent 1.10.1.2:80
   ```

6. Use these commands to display the statistics of SLB pool members and SLB VIPs.

   ```
   show slb stats pool
   show slb stats pool testpool
   show slb stats vip site1
   ```

```
show slb stats vip site2
```

7. To configure the ratio and priority of an existing pool member and to display the current SLB pool statistics, use this command for each pool member, filling in the ipaddress, port, ratio and priority as needed:

   ```
   configure slb pool <poolname> member
   <ipaddress: port> [ratio <ratio> | priority
   <priority>]
   ```

8. Enable SLB and configure it for the appropriate VLANs (client connections enter from the VLAN ***inside***), using these commands:

   ```
   enable slb
   configure inside slb client
   configure server slb server
   ```

9. Enable the routing protocol of choice (in this example, OSPF) and configure it appropriately, using this command:

   ```
   enable ospf
   ```

   See Chapter 13, *RIP and OSPF* for more information.

10. Enable the ESRP protocol on the VLAN ***server*** and configure the ESRP direct-attached hosts mode to allow the proper failover of services, using these commands:

    ```
    enable esrp server
    configure esrp port-mode host ports 1-4, 8
    ```

    The interconnection between the switches is also configured as a host port.

11. Configure SLB to use the ESRP protocol, using this command:

    ```
    configure slb esrp server add unit 1
    ```

Combined SLB and ESRP failover

You can combine SLB and ESRP to provide a high availability topology. Use these two commands to map an ESRP configured VLAN to the SLB failover unit number and to display the current SLB/ESRP configuration:

```
configure slb esrp vlan <vlan name> [add | delete]
unit [1 - 16]
```

```
show slb esrp
```

Configuration of SLB with ESRP

Note the following about the configurations for switches running SLB and ESRP:

- All switch ports connected directly to the servers must be configured as ESRP host ports.
- The link between the two switches must be configured as an ESRP host port.
- The configuration uses transparent mode and HTTP services, but can be configured to support any of the currently supported load balancing protocols.
- Unlike the High Availability configuration, both switches are configured as Switch 1.

Web-Server Configuration

In Figure 18.5, basic HTTP, configured at TCP port 80, is the only service being load balanced. The services must match those configured on the switch. For example, HTTP services configured at TCP port 7080 on the switch require that servers be able to allow connections at port 7080.

Ensure that the SLB connection is valid before trying to transfer the configuration to an ESRP/SLB configuration.

Using High Availability System Features

The 480T routing switch supports several advanced redundant system features. These provide additional assurance that your content is available if a switch experiences a problem. Options include:

- Redundant SLB
- Ping-check
- Active-active operation
- Manual fail-back
- SLB high availability

Redundant SLB

The 480T routing switch supports a failover process that uses a redundant configuration of two switches. If one switch fails, the second switch takes over the SLB duties of the first. By preparing a redundant switch for the possibility of failover, you effectively maintain your site's reliability and availability in advance.

You can configure the switches so that both perform SLB simultaneously. This type of operation is called active-active.

To configure failover, use this command:

```
configure slb failover unit [1 | 2] remote-ip
<ipaddress> local-ip <ipaddress>:<L4Port> {alive-
frequency <seconds> timeout <seconds>} {dead-
frequency <seconds>}
```

```
enable slb failover
```

The switches in a redundant SLB configuration should have identical SLB configurations except for the `failover` parameters. You can configure SLB on one switch, upload the configuration, edit it, and download it to the second switch to replicate the configuration.

Using Ping-Check

Failover ping-check is used to determine if the currently active SLB server has the required network connectivity. If the specified IP address is unreachable for a specified duration, the ping-check triggers a failover to the redundant switch.

The address being pinged should be for a device other than the redundant SLB switch.

To configure ping-check, use these commands:

```
configure slb failover ping-check <ipaddress>
```

```
enable slb failover ping-check
```

Configuring Active-Active Operation

Using active-active redundant SLB, configure one switch as unit 1 and the other switch as unit 2. You then assign the VIPs either to unit 1 or to unit 2 (by default, a VIP is assigned to unit 1).

When both switches are active, each switch performs SLB only for the VIPs assigned to it. If a switch fails, the other switch takes over the VIPs assigned to the failed switch.

The basic failover configure command assigns the switch’s unit number:

```
configure slb failover unit [1 | 2] remote-ip
<ipaddress> local-ip <ipaddress>:<L4Port> {alive-
frequency <seconds> timeout <seconds>} {dead-
frequency <seconds>}
```

where:

- **`remote-ip`**—Specifies the IP address of the redundant SLB switch.
- **`local-ip`**—Specifies the IP address of the switch you are configuring.

All VIPs with a given virtual IP address must be assigned to the same unit.

To assign a VIP to a unit, use this command:

```
configure slb vip <vipname> unit {1 | 2}
```

Sample Active-Active Configuration

Figure 18.6 shows an example of an active-active failover configuration.

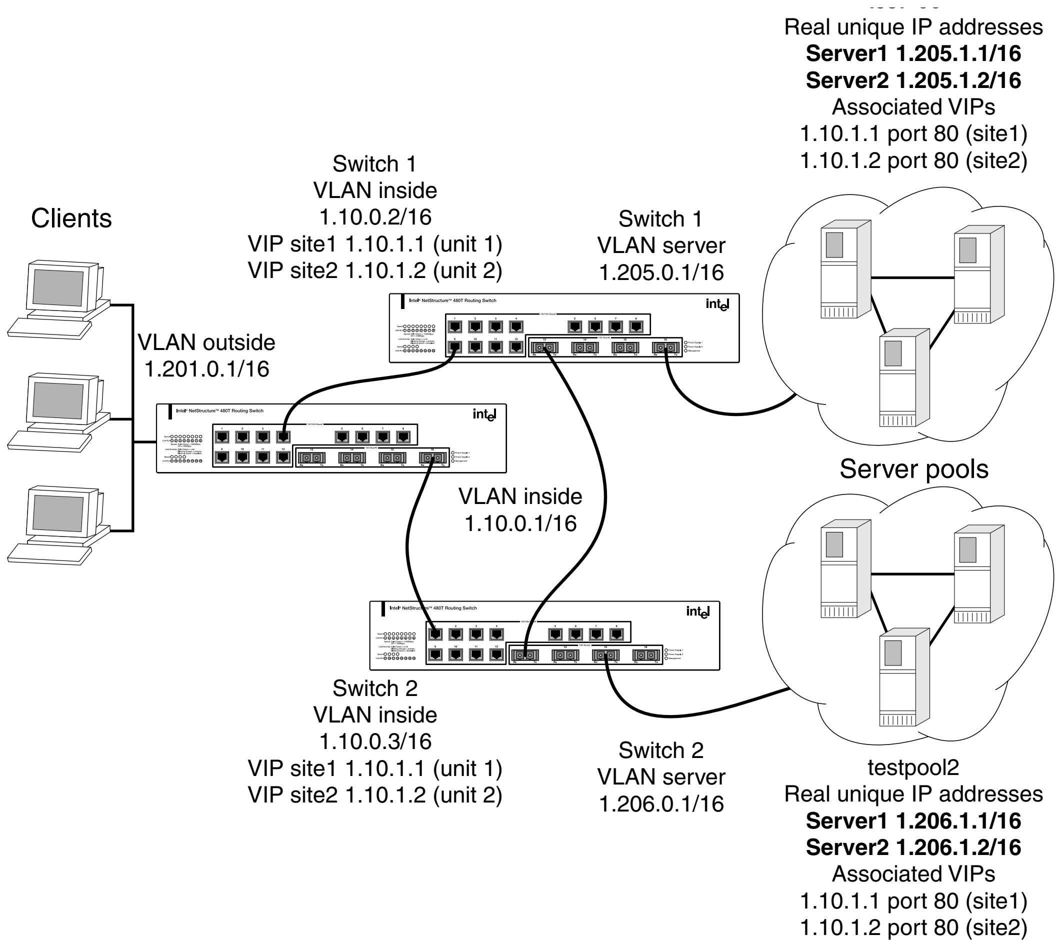

Figure 18.6: Active-active configuration

In this sample configuration, failover is enabled to ensure fault tolerance. To configure this example on the first switch, use these commands:

```
create vlan inside
create vlan server
configure vlan inside ipaddress 1.10.0.2 /16
configure vlan inside add port 10
configure vlan server ipaddress 1.205.0.1 /16
configure vlan server add port 4-8

enable ipforwarding
```

```
create slb pool testpool1
configure slb pool testpool1 add 1.205.1.1:80
configure slb pool testpool1 add 1.205.1.2:80
create slb vip site1 pool testpool1 mode
transparent 1.10.1.1:80
create slb vip site2 pool testpool1 mode
transparent 1.10.1.2:80

configure enable slb
configure vlan inside slb-type client
configure vlan server slb-type server

configure slb failover unit 1 remote 1.10.0.3
local 1.10.0.2:1028

enable slb failover

enable slb failover ping

configure slb vip site1 unit 1
configure slb vip site2 unit 2

configure slb fail ping-check 1.10.0.1 freq 1
```

To configure this example on the second switch, use these commands:

```
create vlan inside
create vlan server
configure vlan inside ipaddress 1.10.0.3 /16
configure vlan inside add port 10
configure vlan server ipaddress 1.206.0.1 /16
configure vlan server add port 4-8

enable ipforwarding

create slb pool testpool2
configure slb pool testpool2 add 1.206.1.1:80
configure slb pool testpool2 add 1.206.1.2:80
create slb vip site1 pool testpool2 mode
transparent 1.10.1.1:80
create slb vip site2 pool testpool2 mode
transparent 1.10.1.2:80

enable slb

configure vlan inside slb-type client

configure vlan server slb-type server

configure slb failover unit 2 remote 1.10.0.2
local 1.10.0.3:1028
```

```
enable slb failover
enable slb fail ping
```

```
configure slb vip site1 unit 1
configure slb vip site2 unit 2
```

```
configure slb fail ping-check 1.10.0.1 freq 1
```

The differences between the configurations of these two switches are the IP addresses, and the designation of the first switch as the master of the active-active configuration.

Using Manual Fail-Back

In an active-active configuration, fail-back is the action of releasing the virtual servers that are assigned to a failed switch when that switch becomes operational again. By default, fail-back occurs automatically. If the minor disruption of fail-back makes automatic fail-back undesirable, you can enable manual fail-back. With manual fail-back, fail-back occurs only when the operator enters the fail-back command.

To enable manual fail-back, use this command:

```
enable slb failover manual-failback
```

To execute a manual fail-back, use this command:

```
configure slb failover failback-now
```

Using SLB High Availability

Using SLB High Availability (SLB H/A) provides redundancy in the case of an SLB service failure. Using SLB H/A, a site is configured with multiple servers spanning two switches. All servers are capable of responding to requests for content, but only those servers connected to the active switch receive requests. The other servers are idle or are used to serve another site.

Figure 18.7 shows an SLB failover configuration using SLB H/A.

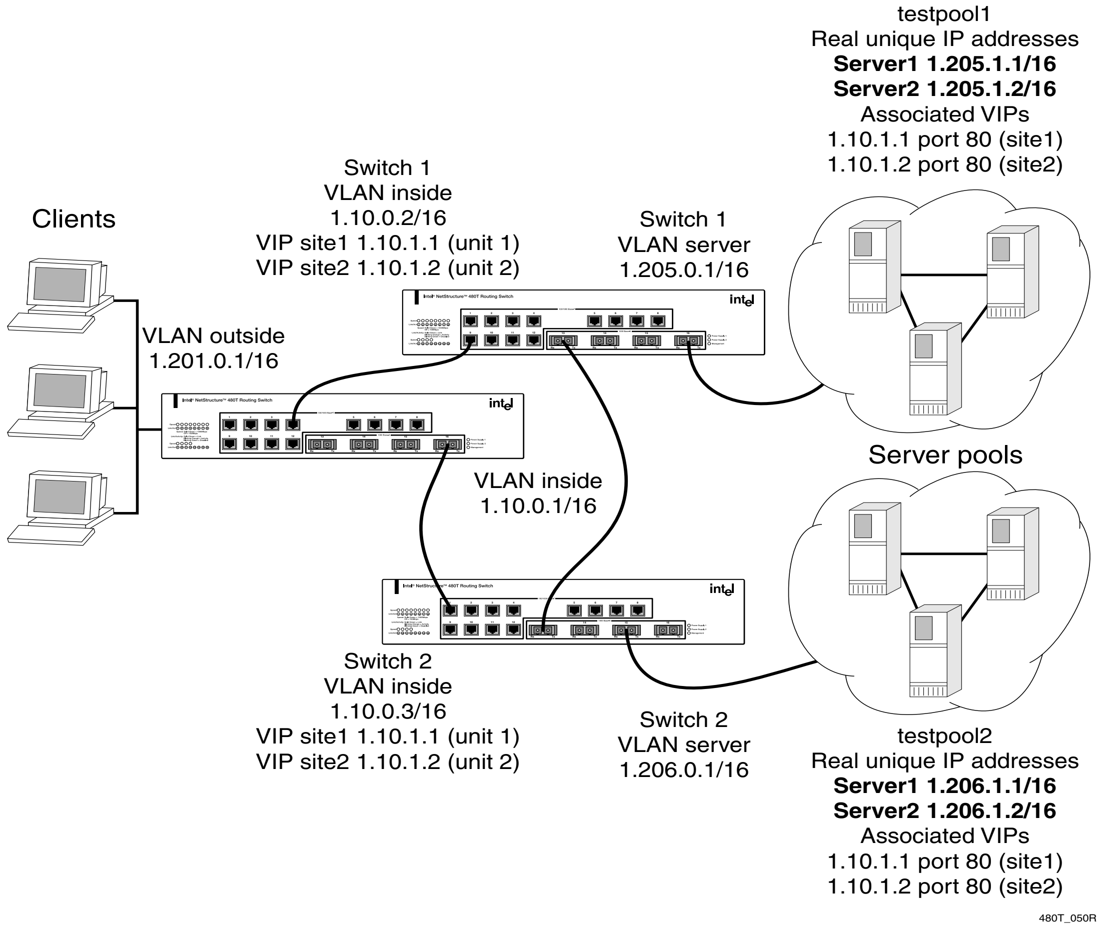

Figure 18.7: SLB failover configuration using SLB H/A

Configuring Clients

The configuration used to connect clients to SLB virtual sites with High Availability enabled is transparent to the accessing clients. As with normal SLB, the clients connect to the VIP believing that it is the physical address on a host server.

Configuring Switches for SLB H/A

The procedure used to configure the two switches for SLB High Availability is described below.

Create the VLANs, using these commands:

```
create vlan inside
```

```
create vlan server
```

The VLAN ***inside*** connects to the gateway and the VLAN ***server*** contains all of the load balanced servers.

The gateway is connected to the VLAN ***inside***, using these commands:

```
configure inside ipaddress 1.10.0.2 /16
```

```
configure inside add port 10
```

Connect the servers to the VLAN ***server*** on ports 4-8, using these commands:

```
configure server ipaddress 1.205.0.1 /16
```

```
configure server add port 4-8
```

Two servers are connected to each High Availability switch.

Enable IP forwarding, create a server pool called ***testpool1***, and add two servers to ***testpool1*** using TCP port 80, using these commands:

```
enable ipforwarding
```

```
create slb pool testpool1
```

```
configure slb pool testpool1 add 1.205.1.1:80
```

```
configure slb pool testpool1 add 1.205.1.2:80
```

Create SLB VIP addresses for the two Web sites (***site1*** and ***site2***) and associate the server pool ***testpool*** with them, using these commands:

```
create slb vip site1 pool testpool1 mode
transparent 1.10.1.1:80
```

```
create slb vip site2 pool testpool1 mode
transparent 1.10.1.2:80
```

Then create ***testpool2*** and add 1.206.1.1:80 and 1.206.1.2:80 to it. Create an identical SLB for ***testpool2.***

Then, enable SLB and configure it for the appropriate VLANs (client connections enter from the VLAN ***inside***), using these commands:

```
enable slb
```

```
configure inside slb client
```

```
configure server slb server
```

Configure SLB H/A for the switch, using this command:

```
configure slb failover unit 1 remote 1.10.0.3
local 1.10.0.2 l4-port 1028
```

One switch in a High Availability pair is designated as unit 1 and the other is designated as unit 2. VIPs associated with the unit numbers are primarily serviced by the appropriate switch. The IP address of the remote switch in the failover pair is 1.10.0.3. The IP address of the local interface used by the High Availability protocol to communicate with the remote switch is 1.10.0.2. The Layer 4 port used by the High Availability protocol to exchange information is 1028.

Along with performing normal status checking on the remote switch, the High Availability protocol pings the gateway to ensure that a connection to the client exists. If the connection to the gateway at IP address 1.10.0.1 fails, the remote switch services all of the connections. Configure status checking and enable failover using these commands:

```
enable slb failover
```

```
configure slb failover ping-check 1.10.0.1
```

```
enable slb failover ping
```

Configure the unit numbers on the two sites to determine which of the High Availability switches will actively serve the VIPs, using these commands:

```
configure slb vip site1 unit 1
```

```
configure slb vip site2 unit 2
```

In this example, ***site1*** is serviced by the current switch and the remote switch (configured as unit 2) services ***site2***. A switch configured as unit 1 services unit 2's VIPs only when the remote switch (configured as unit 2) fails.

Notes on Configuring SLB H/A

These are important notes about the configurations for SLB H/A:

- In the design shown in Figure 18.7, only the servers directly connected to the switch that is actively servicing the VIP are used in the load balancing scheme. Without ESRP, another switch interconnecting all the servers is necessary.

- One switch is designated as unit 1 and the other as unit 2. This designation determines which VIPs are active on each switch in the failover pair.
- In this configuration, ***site1*** is serviced by Switch 1 and has two servers that respond to client requests. ***Site2*** is serviced by the remote switch (Switch 2), and has two other servers that respond to client requests.
- If ping-check is enabled, it must not be directed at the remote switch. The remote switch is checked by the High Availability protocol. The ping-check works best when directed at a gateway to ensure that a path out of the network is available to the switch.
- The configuration uses transparent mode and HTTP services, but can be configured to support any of the currently supported load balancing protocols.
- The configurations for the High Availability switches are identical, with the exception of the **`failover`** command:

  ```
  configure slb failover unit 1 remote 1.10.0.3
  local 1.10.0.2 l4-port 1028
  ```

- The remote switch is set to unit 2, and the remote/local IP addresses are reversed to accurately describe the network, as shown in this command:

  ```
  configure slb failover unit 2 remote 1.10.0.2
  local 1.10.0.3 l4-port 1028
  ```

Web Server configuration

In the configuration shown in Figure 18.7 on page 388, basic HTTP, configured at TCP port 80, is the only service being load balanced. It is important that the services match those configured on the switch. For example, HTTP services configured at TCP port 7080 on the switch would require the servers to be able to allow connections at port 7080. You must also ensure that the SLB configuration is valid before enabling High Availability.

All four servers (two local and two connected to the remote switch) should be identical in content, with the content for both ***site1*** and ***site2*** configured to be served.

This configuration uses transparent mode. Therefore, the VIPs need to be added to the servers as loopback addresses. This is done by configuring the network interfaces on the servers. A detailed description for doing this is provided after Figure 18.1.

Advanced SLB Commands

Table 18.4 describes advanced SLB commands. For further command options, press the Tab key in the command line interface.

Table 18.4: Advanced SLB Commands

Command	Description
clear slb persistence {vip <vip name>}	Resets all connection information in the persistence table. New connections opened are directed to a new server.
clear slb connections {ip address <ipaddress>: L4Port | vip <vip name>}	Resets all connections.
configure slb failover failback-now	Configures the local SLB to release remote SLB resources if the remote SLB is alive.
configure slb failover ping-check <ipaddress> {frequency <seconds> timeout <seconds>}	Configures the SLB device to actively determine if an external gateway is reachable by performing a ping. If the external gateway is not reachable, the VIPs failover to the remote SLB device. Specify: • **ipaddress**—The IP address of the external gateway. • **frequency**—The interval, in seconds, between pings sent to the remote gateway. The default setting is 1. • **timeout**—The amount of time, in seconds, before the local device declares the remote gateway is not reachable. The default setting is 3.

Table 18.4: Advanced SLB Commands

Command	Description
configure slb failover unit <number> {remote-ip <ipaddress> local-ip <ipaddress>:{<L4Port>}}	Configures the slb failover. Specify: • **`remote-ip-address`**—The remote peer IP address. • **`local-ip-address`**—The address of a local IP interface used for the failover connection. • **`L4Port`**—The TCP port used for keep-alives between the failover peers. The default port is 1028. • **`unit`**—The unit number for this SLB device. The default unit number is 1.
configure slb global [ping-check | tcp-port-check | service-check] frequency <seconds> timeout <seconds>	Configures default health checking frequency and timeout period. If the health check frequency and timeout are not specified for a specific node or VIP, the global values are used. Specify one of these service checkers: • **`ping-check`**—L3-based pinging of the physical node. Default ping frequency is one ping generated to the node each 10 seconds. If the node does not respond to any ping within a timeout period of 30 seconds (3 ping intervals), the node is considered inoperable. • **`tcp-port-check`**—L4-based TCP port open/close testing. Default values are 30 seconds for frequency and 90 seconds for timeout. • **`service-check`**—L7-based application-dependent checking. Default values are 60 seconds for frequency and 180 seconds for timeout.
configure slb global ftp userid <userid> password {encrypted} {<password>}	Configures default parameters for L7 service checking. If no password is provided, you are prompted twice.

Table 18.4: Advanced SLB Commands

Command	Description
configure slb global http url <url_string> match-string [<match_string> | any-content]	Configures the default parameters for L7 service checking.
configure slb global nntp newsgroup <newsgroup>	Configures the default parameter for L7 service checking.
configure slb global persistence-level [same-vip-same-port | same-vip-any-port | any-vip]	Configures the default parameter for persistence level.
configure slb global persistence-method [per-packet | per-session]	Configures the default parameter for persistence method.
configure slb global pop3 userid <userid> password {encrypted} {<password>}	Configures the default parameter for L7 service checking.
configure slb global smtp <dns_domain>	Configures the default parameter for L7 service checking.
configure slb global synguard max-unacknowledge-SYNs <num_syns>	Configures the `num_syns` value that is used to trigger the SYN-guard feature.
configure slb global telnet userid <userid> password {encrypted} {<password>}	Configures default parameters for L7 service checking. If no password is provided, you are prompted twice for the password.
configure slb node <ipaddress>:<L4Port> tcp-port-check frequency <seconds> | timeout <seconds>	Overrides the global default frequency and timeout values for this node. Use a value of 0 to restore settings to global default values.
configure slb node <ipaddress> ping-check frequency <seconds> timeout <seconds>]	Overrides the global default frequency and timeout values for this node. Use a value of 0 to restore the settings to the global default values.
configure slb vip <vipname> max-connections <connections>	Configures the maximum connections allowed to a particular VIP. A value of 0 indicates that no maximum is enforced. The default value is 0.

Table 18.4: Advanced SLB Commands

Command	Description
configure slb vip <vipname> service-check frequency <seconds> timeout <seconds>	Configures the L7 service check frequency and timeout parameters for a particular VIP. To return to the global values, specify `0` for frequency and timeout.
configure slb vip <vipname> service-check http {url <url> match-string [<match_string> | any-content]}	Configures VIP service checking for the HTTP service. When the match-string option is specified, the string must be in the first 1000 bytes of the returned Web page.
configure slb vip <vipname> service-check ftp {userid <userid> | password {encrypted} <password>}	Configures VIP service checking for the FTP service.
configure slb vip <vipname> service-check telnet {userid <userid> | password {encrypted} <password>}	Configures VIP service checking for the telnet service.
configure slb vip <vipname> service-check smtp {<dns_domain>}	Configures VIP service checking for the SMTP service.
configure slb vip <vipname> service-check nntp <newsgroup>	Configures VIP service checking for the NNTP service.
configure slb vip <vipname> service-check pop3 userid <userid> password {encrypted} {<password>}	Configures VIP service checking for the POP3 service.
configure slb vip <vipname> unit <number>	Configures a unit number of a VIP name for active-active failover. The default unit number is 1.
disable slb failover	Disables SLB failover.
disable slb failover manual-failback	Disables manual failback.
disable slb failover ping-check	Disables ping-check to an external gateway.
disable slb global synguard	Disables the TCP SYN-guard feature.

Table 18.4: Advanced SLB Commands

Command	Description
disable slb node <ipaddress>:{<L4Port> | all} tcp-port-check	Disables L4 port checking.
disable slb node <ipaddress> ping-check	Disables L3 pinging.
disable slb vip [<vipname> | all] client-persistence	Disables client-persistence.
disable slb vip [<vipname> | all] close-connections-now	Disables one or all VIP groups. All existing connections are immediately closed.
disable slb vip [<vipname> | all] service-check	Disables L7 service checking.
disable slb vip [<vipname> | all] sticky-persistence	Disables sticky persistence.
disable slb vip [<vipname> | all] svcdown-reset	Disables svcdown-reset.
enable slb failover	Enables the SLB failover mechanism. The default setting is disabled.
enable slb failover manual-failback	Enables manual failback.
enable slb failover ping-check	Enables ping-checking to an external gateway. The default setting is disabled.
enable slb global synguard	Enables the TCP SYN-guard feature. The SYN-guard feature minimizes the effect of the TCP-open type of denial-of-service attack by keeping track of all the half-open connections. When the number of half-open connections exceeds the **num_syns** value, the half-open connections are fast-aged out.
enable slb node <ipaddress> ping-check	Enables L3 pinging to the node address. Ping-check is automatically enabled when a node is added to a pool.
enable slb node <ipaddress>:<L4Port> tcp-port-check	Enables L4 port-check to the node address.

Table 18.4: Advanced SLB Commands

Command	Description
enable slb vip [<vipname> | all] client-persistence {mask <mask>}	Enables client persistence and specifies the timeout and client address mask. If the client sets up multiple sessions to a virtual server, all sessions must connect to the same physical node. Enabling client persistence instructs the switch to forward new session requests from the same client (or clients on the same network using the **mask** argument) to the same node. The association between the client and physical node is ended after the specified timeout. The default is disabled.
enable slb vip [<vipname> | all] service-check	Enables L7 service checking based on: • If a service check is already configured, it uses the user-configured service-checking information. • If a service-check is not explicitly configured and a well-known port is used when creating a VIP, the switch guesses the application based on the well-known port number and starts the L7 service checker with the global default parameters.
enable slb vip [<vipname> | all] sticky-persistence {netmask <mask>}	Enables sticky persistence and specifies the timeout. Sticky persistence is usually used to load balance firewall and Web caches. When enabled, the switch forwards all traffic and new sessions toward a destination address (or address within a certain subnet boundary specified by the **mask** argument) to the same physical node. The default setting is disabled.

Table 18.4: Advanced SLB Commands

Command	Description
enable slb vip [<vipname> | all] svcdown-reset	Enables the svcdown-reset configuration. If enabled, the switch sends TCP RST to both the clients and the node, if the node associated with this VIP completely fails a ping-check, port-check, or service-check. Otherwise, the connections to the node are left as is, and are subject to connection reaping if idle for longer than the treaper-timeout configured on the SLB port. The default setting is disabled.
show slb failover	Disables the SLB failover configuration and status.
unconfigure slb vip [<vipname> | all] service-check	Disables and removes the service check configuration.

Web Cache Redirection

Web cache redirection uses the TCP or UDP port number to redirect client requests to a target device (or group of devices). Web cache redirection transparently redirects traffic to Web cache devices or to proxy servers and firewalls located in a demilitarized zone.

There are two ways to configure Web cache redirection:

- Transparent mode SLB (described earlier in this chapter)
- Flow redirection

Flow Redirection

Flow redirection examines traffic and redirects it based on these criteria:

- IP source address and mask
- IP destination address and mask
- Layer 4 port

Precedence of Flow Redirection Rules

Multiple flow redirection rules can overlap in making a redirection decision. In these cases, precedence is determined by "best match" where the most specific redirection rule that satisfies the criteria will win. The best match is determined in this order:

- Destination IP Address/Mask
- Destination IP Port
- Source IP Address/Mask

In general, these rules apply:

- If a flow with a comparatively better matching mask on an IP address satisfies the content of a packet, that flow is observed.
- If one flow redirection rule contains any as an L4 protocol and a second flow redirection rule contains explicit L4 port information, the second is observed, if the packet contains matching L4 information.
- If one flow has a comparatively better match on source information and a second flow has comparatively better match on destination information then the rule with the better match on the destination information is selected.

For example, in the following two cases, the rule with the best match is the rule that is selected.

Table 18.5: Example #1: Flow Redirection Rules

Destination IP Address	Destination IP Port	Source IP Address	Priority Selection
192.0.0.0/8	80	ANY	1
192.168.0.0/16	ANY	ANY	2

In this case, Policy 1 is the rule with the best match as it contains an explicit Destination IP Port even though the mask for the Destination IP Address is less specific.

Table 18.6: Example #2: Flow Redirection Rules

Destination IP Address	Destionation IP Port	Source IP Address	Priority Selection
192.168.2.0/24	**80**	**ANY**	**2**
192.168.0.0/16	**ANY**	**10.10.10.0./24**	**4**
192.168.2.0/24	**ANY**	**10.10.0.0/16**	**3**
192.168.2.0/24	**80**	**10.10.0.0/16**	**1**

In this case, Policy 4 is the rule with the best match as it contains an explicit destination IP Port.

Flow Redirection Commands

To configure flow redirection, use the commands listed in Table 18.7. For further command options, press the Tab key in the command line interface.

Table 18.7: Flow Redirection Commands

Command	Description
configure flow-redirect <flow_policy> add next-hop <ipaddress>	Adds the next-hop host (gateway) that is to receive the packets that match the flow policy. By default, ping-based health checking is enabled.
configure flow-redirect <flow_policy> delete next-hop <ipaddress>	Deletes the next-hop host (gateway).
configure flow-redirect <flow-policy> service-check [ftp | http | L4-port | nntp | ping | pop3 | smtp | telnet]	Adds a service check for the specified service to the flow redirection policy

Table 18.7: Flow Redirection Commands (continued)

Command	Description
create flow-redirect <flow_policy> [any | tcp | udp] destination [<ipaddress/mask> | any] ip-port [<L4Port> | any] source [<ipaddress/mask> | any]	Creates a flow redirection policy.
delete flow-redirect <flow_policy>	Deletes a flow redirection policy.
show flow-redirect	Displays the current flow redirection configuration and statistics.

Flow Redirection Example

Figure 18.8 uses flow redirection to redirect Web traffic to Web cache servers. In this example, the clients and the cache devices are located on different networks. This is done by creating a different VLAN for the clients and cache devices.

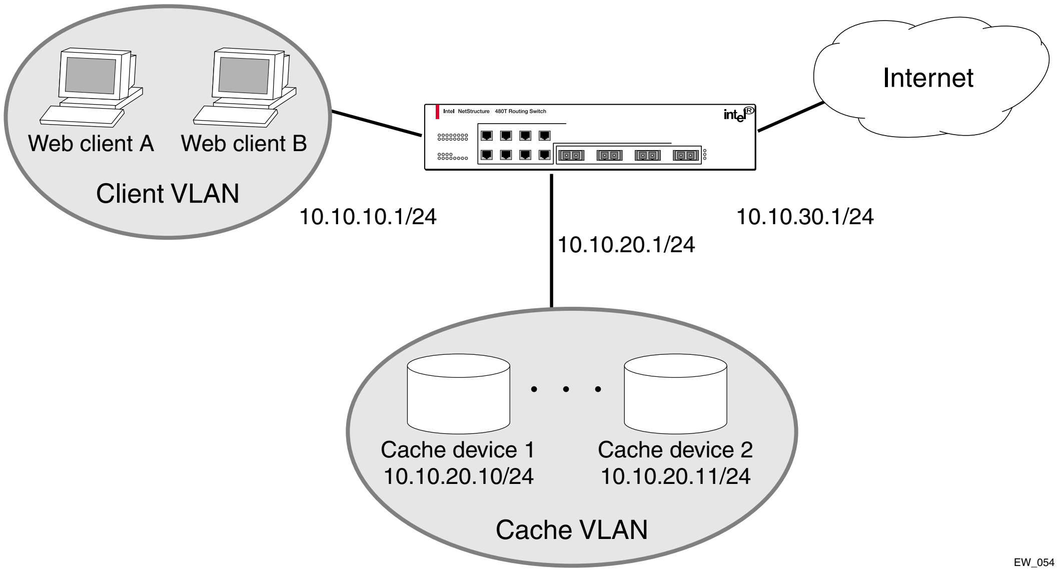

Figure 18.8: Flow-redirection example

These commands are used to configure the 480T routing switch in this example:

```
create vlan client
configure vlan client add port 1
configure vlan client ipaddress 10.10.10.1/24

create vlan cache
configure vlan cache add port 2
configure vlan cache ipaddress 10.10.20.1/24

create vlan internet
configure vlan internet add port 3
configure vlan internet ipaddress 10.10.30.1/24

enable ipforwarding

create flow-redirect wcr tcp destination any ip-
port 80 source any
configure flow-redirect wcr add next-hop
10.10.20.10
configure flow-redirect wcr add next-hop 10.10.20.1
```

19 Status Monitoring and Statistics

This chapter describes how to view the current operating status of the Intel® NetStructure™ 480T routing switch, how to display information in the log, and how to take advantage of available Remote Monitoring (RMON) capabilities.

Viewing statistics on a regular basis allows you to see how well your network is performing. If you keep simple daily records, you may see trends emerging and notice problems arising before they cause major network faults.

Status Monitoring

The status monitoring facility provides information about the 480T routing switch. This information may be useful when contacting Intel Customer Support, should you have a problem. The local management software includes many **`show`** commands that display information about different switch functions and facilities.

Table 19.1 describes **`show`** commands that are used to monitor the status of the 480T routing switch . For further command options, press the Tab key in the command line interface.

Table 19.1: Status Monitoring Commands

Command	Description
show log config	Displays the log configuration, including the syslog host IP address, the priority level of messages being logged locally, and the priority level of messages being sent to the syslog host.
show log {<priority>}	Displays the current snapshot of the log. Priority options filter the log to display messages with the selected priority or higher (more critical). Specify: • **`Critical`** • **`Emergency`** • **`Alert`** • **`Error`** • **`Warning`** • **`Notice`** • **`Info`** • **`Debug`** • **`Configuration`** If not specified, all messages display.
show memory {detail}	Displays the current system-memory information. Specify the **`detail`** option to view task-specific memory usage.

Table 19.1: Status Monitoring Commands (continued)

Command	Description
show switch	Displays the current switch information, including: • sysName, sysLocation, sysContact • MAC address • Current time and date, system uptime, and time zone • Operating environment (temperature indication, fans, and power supply status) • Non-Volatile Random Access Memory (NVRAM) configuration information • Scheduled reboot information • Software licensing information
show version	Displays the hardware and software versions running on the switch.

Port Statistics

The 480T routing switch allows you to view port statistic information. The summary information lists values for the current counter against each port on each operational module in the system, and it is refreshed approximately every 2 seconds. Values are displayed to nine digits of accuracy.

To view port statistics, use this command:

```
show ports <portlist> stats
```

This port statistic information is collected:

- **Link Status**—The current status of the link. Options are:
  - Ready: the port is ready to accept a link
  - Active: the link is present at this port
- **Transmitted Packet Count (Tx Pkt Count)**—The number of packets that were successfully transmitted by the port.
- **Transmitted Byte Count (Tx Byte Count)**—The total number of data bytes successfully transmitted by the port.

- **Received Packet Count (Rx Pkt Count)**—The total number of good packets that were received by the port.
- **Received Byte Count (Rx Byte Count)**—The total number of bytes that were received by the port, including bad or lost frames. This number includes bytes contained in the Frame Check Sequence (FCS), but excludes bytes in the preamble.
- **Received Broadcast (Rx Bcast**)—The total number of frames received by the port that are addressed to a broadcast address.
- **Received Multicast (Rx Mcast)**—The total number of frames received by the port that are addressed to a multicast address.

Port Errors

The 480T routing switch tracks errors for each port. To view port transmit errors, use this command:

```
show ports <portlist> txerrors
```

This port transmit error information is collected:

- **Port Number**
- **Link Status**—The current status of the link. Options are:
  - Ready: the port is ready to accept a link
  - Active: the link is present at this port
- **Transmit Collisions (Tx Coll)**—The total number of collisions seen by the port, regardless of whether a device connected to the port participated in any of the collisions.
- **Transmit Late Collisions (Tx Late Coll)**—The total number of collisions that occurred after the port's transmit window expired.
- **Transmit Deferred Frames (Tx Deferred)**—The total number of frames that were transmitted by the port after the first transmission attempt was deferred by other network traffic.
- **Transmit Errored Frames (Tx Error)**—The total number of frames that were not completely transmitted by the port because of network errors (such as late collisions or excessive collisions).
- **Transmit Parity Frames (Tx Parity)**—The bit summation has a parity mismatch.

To view port receive errors, use this command:

```
show ports <portlist> rxerrors
```

The following port receive error information is collected:

- **Receive Bad CRC Frames (Rx CRC)**—The total number of frames received by the port that were of the correct length, but contained a bad FCS value.
- **Receive Oversize Frames (Rx Over)**—The total number of good frames the port received that were longer than the supported maximum length of 1,522 bytes. Ports with jumbo frames enabled do not increment this counter.
- **Receive Undersize Frames (Rx Under)**—The number of frames the port received that were less than 64 bytes long.
- **Receive Fragmented Frames (Rx Frag)**—The total number of frames the port received that were of incorrect length and contained a bad FCS value.
- **Receive Jabber Frames (Rx Jab)**—The total number of frames the port received, greater than the support maximum length and that had a Cyclic Redundancy Check (CRC) error.
- **Receive Alignment Errors (Rx Align)**—The total number of frames received by the port that occurs if a frame has a CRC error and does not contain an integral number of octets.
- **Receive Frames Lost (Rx Lost)**—The total number of frames that were received by the port that were lost due to buffer overflow in the switch.

Port Monitoring Display Keys

Table 19.2 describes the keys used to control the displays that appear when you issue any of the `show port` commands.

Table 19.2: Port Monitoring Display Keys

Key(s)	Description
U	Displays the previous page of ports.
D	Displays the next page of ports.

Table 19.2: Port Monitoring Display Keys (continued)

Key(s)	Description
Esc or Enter	Exits from the screen.
0	Clears all counters.
Spacebar	Cycles through these screens: • Packets per second • Bytes per second • Percentage of bandwidth Available using the **`show port utilization`** command only.

Setting the System Recovery Level

You can configure the system to automatically reboot after a software task exception, using this command:

```
configure sys-recovery-level [none | critical |
all]
```

Where:

- **`none`**—Configures the level to no recovery.
- **`critical`**—Configures the switch to log an error into the syslog and automatically reboot the system after a critical task exception.
- **`all`**—Configures the switch to log an error into the syslog and automatically reboot the system after any task exception.

The default setting is **`none`**.

Logging

The 480T routing switch log tracks all configuration and fault information pertaining to the device. Each entry in the log contains this information:

- **Timestamp** The timestamp records the month and day of the event, along with the time (hours, minutes, and seconds) in the form HH:MM:SS. If the event was caused by a user, the user name is also provided.

- **Fault level**—Table 19.3 describes the three levels of importance that the system can assign to a fault.

Table 19.3: Fault Levels

Level	Description
Critical	A desired switch function is inoperable. The switch may need to be reset.
Warning	A noncritical error that may lead to a function failure.
Informational	Actions and events that are consistent with expected behavior.
Debug	Information that is useful when performing detailed troubleshooting procedures.

By default, log entries that are assigned a critical or warning level remain in the log after a switch reboot. Issuing a clear log command does not remove these static entries. To remove log entries of all levels (including warning or critical), use this command:

```
clear log static
```

- **Subsystem**—The subsystem refers to the specific functional area to which the error refers. Table 19.4 describes the subsystems

Table 19.4: Fault Log Subsystems

Subsystem	Description
Syst	General system-related information. Examples include memory, power supply, security violations, fan failure, overheat condition, and configuration mode.
STP	Shielded Twisted Pair (STP) information. Examples include an STP state change.

Table 19.4: Fault Log Subsystems (continued)

Subsystem	Description
Brdg	Bridge-related functionality. Examples include low table space and queue overflow.
SNMP	SNMP information. Examples include community string violations.
Telnet	Information related to Telnet login and configuration performed using a Telnet session.
VLAN	VLAN-related configuration information.
Port	Port management-related configuration. Examples include port statistics and errors.

Local Logging

The 480T routing switch maintains 1,000 messages in its internal log. You can display a snapshot of the log at any time, using the command:

```
show log {<priority>}
```

Displays the current snapshot of the log. Priority filters the log to display messages with the selected or higher (more critical) priority. Priorities include (in order):

- Critical
- Emergency
- Alert
- Error
- Warning
- Notice
- Info
- Debug

If not specified, info and higher priority messages display.

Real-Time Display

Along with viewing a snapshot of the log, you can configure the system to maintain a running real-time display of log messages on the console. To turn on the log display, enter this command:

```
enable log display
```

To configure the log display, use this command:

```
configure log display {<priority>}
```

If `priority` is not specified, only messages of critical priority display.

If you enable the log display on a terminal connected to the console port, your settings remain in effect even after your console session ends (unless you explicitly disable the log display).

When using a Telnet connection, if your Telnet session is disconnected (because of the inactivity timer or for other reasons), the log display is automatically halted. To restart the log display, use the `enable log display` command.

Remote Logging

Along with maintaining an internal log, the 480T routing switch supports remote logging using the syslog host facility. You can configure up to four syslog servers for remote logging. To enable remote logging configure the syslog host to accept and log messages. Use these commands:

1. Enable remote logging by using this command:

   ```
   enable syslog
   ```

2. Configure remote logging by using this command:

   ```
   configure syslog {add} <ipaddress> <facility>
   {<priority>}
   ```

3. Specify:

   - `ipaddress`—The IP address of the syslog host.
   - `facility`—The syslog facility level for local use. Options include `local0` through `local7`.
   - `priority`—Filters the log to display messages with the selected or higher (more critical) priority.

The priorities are the same as for local logging.

If not specified, only critical priority messages are sent to the syslog host.

Logging Configuration Changes

The local management software allows you to record all configuration changes (and their sources) made through the CLI using Telnet or the local console. The changes are logged to the system log.

Each log entry includes the user account name that performed the change and the source IP address of the client (if Telnet was used).

Configuration logging applies only to commands that result in a configuration change. To enable configuration logging, use this command:

```
enable cli-config-logging
```

To disable configuration logging, use this command:

```
disable cli-config-logging
```

CLI configuration logging is enabled by default.

Logging Commands

The commands described in Table 19.5 allow you to configure logging options, reset the options, display the log, and clear the log. For further command options, press the Tab key in the command line interface.

Table 19.5: Logging Commands

Command	Description
clear counters	Clears all switch statistics and port counters.
clear log {static}	Clears the log. If `static` is specified, the critical log messages are also cleared.

Table 19.5: Logging Commands (continued)

Command	Description
configure log display {<priority>}	Configures the real-time log display. Displays the current snapshot of the log. Priority filters the log to display messages with the selected or higher (more critical) priority. Priorities include (in order): • **`Critical`** • **`Emergency`** • **`Error`** • **`Alert`** • **`Warning`** • **`Notice`** • **`Info`** • **`Debug`** If not specified, **`info`** and higher priority messages display.
configure syslog {add} <ip_address> <facility> {<priority>}	Configures the syslog host address and filters messages sent to the syslog host. You can configure up to four syslog servers. Options include: • **`ipaddress`**—The IP address of the syslog host. • **`facility`**—The syslog facility level for local use (local0 - local7). • **`priority`**—The priority filter as described in the previous command, one of critical, emergency, error, alert, warning, notice, info or debug.
configure syslog delete <ip_address> <facility> {<priority>}	Deletes a syslog host address.
disable cli-config-logging	Disables configuration logging.

Table 19.5: Logging Commands (continued)

Command	Description
disable log display	Disables the log display.
disable syslog	Disables logging to a remote syslog host.
enable cli-config-logging	Enables the logging of CLI configuration commands to the Syslog for auditing purposes. The default setting is enabled.
enable log display	Enables the log display.
enable syslog	Enables logging to a remote syslog host.
show log config	Displays the log configuration, including the syslog host IP address, the priority level of messages being logged locally, and the priority level of messages being sent to the syslog host.
show log {<priority>}	Displays the current snapshot of the log. `Priority` filters the log to display messages with the selected or higher (more critical) priority. Priorities is one of Critical, Emergency, Error, Alert, Warning, Notice, Info, or Debug. If not specified, info and higher priority messages display.

RMON

Using the Remote Monitoring (RMON) capabilities of the 480T routing switch allows network administrators to improve system efficiency and reduce the load on the network.

This sections explain more about the RMON concept and the RMON features supported by the switch.

RMON is the common abbreviation for the Remote Monitoring Management Information Base (MIB) system defined by the Internet Engineering Task Force (IETF) documents RFC 1271 and RFC 1757. You can use RMON to monitor LANs remotely.

A typical RMON setup consists of two components:

- **RMON probe**—An intelligent, remotely controlled device or software agent that continually collects statistics about a LAN segment or VLAN. The probe transfers the information to a management workstation on request, or when a predefined threshold is crossed.
- **Management workstation**—Communicates with the RMON probe and collects the statistics from it. The workstation does not have to be on the same network as the probe, and can manage the probe by in-band or out-of-band connections.

You can only use the RMON features of the system if you have an RMON management application, and have enabled RMON on the switch.

RMON Features

The IETF defines nine groups of Ethernet RMON statistics. The 480T routing switch supports four of these groups:

- Statistics
- History
- Alarms
- Events

Statistics

The RMON Ethernet Statistics group provides traffic and error statistics showing packets, bytes, broadcasts, multicasts, and errors on a LAN segment or VLAN.

Information from the Statistics group is used to detect changes in traffic and error patterns in critical areas of the network.

History

The History group provides historical views of network performance by taking periodic samples of the counters supplied by the Statistics group. The History group features user-defined sample intervals and bucket counters for complete customization of trend analysis.

The History group is useful for analysis of traffic patterns and trends on a LAN segment or VLAN, and to establish baseline information indicating normal operating parameters.

Alarms

The Alarms group provides a versatile, general mechanism for setting threshold and sampling intervals to generate events on any RMON variable. Both rising and falling thresholds are supported, and thresholds can be on the absolute value of a variable or its delta value. Alarm thresholds may be auto-calibrated or set manually.

Alarms inform you of a network performance problem and can trigger automated action responses through the Events group.

Events

The Events group creates entries in an event log or sends SNMP traps to the management workstation. An event is triggered by an RMON alarm. You can configure the switch to:

- Ignore the event
- Log the event
- Send an SNMP trap to the receivers listed in the trap receiver table
- Both log and send a trap

The RMON traps are defined in RFC 1757 for rising and falling thresholds.

Effective use of the Events group saves you time. Rather than having to watch real-time graphs for important occurrences, you can depend on the Events group for notification.

Through the SNMP traps, events can trigger other actions, providing a mechanism for an automated response to certain occurrences.

Configuring RMON

RMON requires one probe per LAN segment, and stand-alone RMON probes have traditionally been expensive. Therefore, Intel’s approach is to provide an affordable RMON probe into the agent of each system. This allows RMON to be widely deployed around the network without costing more than traditional network management.

The 480T routing switch accurately maintains RMON statistics at the maximum line rate of all of its ports. For example, statistics can be related to individual ports.

RMON Probe with Security Features Enabled

A probe must be able to monitor all traffic. Unlike Intel's built-in probe, a stand-alone probe must be attached to a nonsecure port. Implementing RMON in the 480T routing switch allows for all ports to have security features enabled.

To enable or disable the collection of RMON statistics on the switch, use this command:

```
[enable | disable] rmon
```

By default, RMON is disabled. However, even in the disabled state, the 480T routing switch responds to RMON queries and sets for alarms and events. By enabling RMON, the switch begins the processes necessary for collecting switch statistics.

Event Actions

The actions that you can define for each alarm are shown in Table 19.6.

Table 19.6: Event Actions

Action	High Threshold
No action	
Notify only	Send trap to all trap receivers.
Notify and log	Send trap; place entry in RMON log.

To be notified of events using SNMP traps, you must configure one or more trap receivers.

20 Software Upgrade and Boot Options

Overview

This chapter describes the procedure for upgrading the Intel® NetStructure™ 480T routing switch firmware image. It also includes a discussion of how to save and load a primary and secondary image and configuration file on the switch.

Saving Configuration Changes

The configuration is the customized set of parameters that you have selected to run on the switch. As you make configuration changes, the new settings are stored in run-time memory. Settings that are stored in run-time memory are not retained by the switch when the switch is rebooted.

To retain the settings and have them load when you reboot the switch, you must save the configuration to non-volatile (more permanent) storage.

The switch can store two different configurations: a primary and a secondary. When you save configuration changes, you can select the configuration you want to save the changes to. If you do not specify, the changes are saved to the configuration area currently in use.

If you make a mistake, or find you must revert to the configuration as it was before you started making changes, you can set the switch to use the secondary configuration on the next reboot.

If the switch is rebooted during a configuration save, the switch boots to factory default settings. The configuration in the process of being saved is unaffected.

To save the configuration, use this command:

```
save {configuration} {primary | secondary}
```

To use the configuration, use this command:

```
use configuration [primary | secondary]
```

The configuration takes effect on the next reboot.

Upgrading Your Switch

To upgrade your 480T routing switch you may need to upgrade the BootROM image ***and*** *the firmware. Refer to the Late Breaking News document at http://support.intel.com.*

The image file contains the factory-installed executable code or program that runs on the switch. As new versions of the image are released, you should upgrade the firmware and the BootROM image running on your switch.

The images are upgraded by using a download procedure from a TFTP (Trivial File Transfer Protocol) server on the network your switch is connected to.

To upgrade the switch, you must:

1. Save your configuration to the TFTP server.
2. Download the new BootROM and reboot your switch.
3. Download the new firmware and reboot your switch.
4. Restore your configuration from the TFTP server.

Since the switch stores both a primary and a secondary configuration, you can upgrade the firmware into the primary configuration, while retaining the older versions in the secondary configuration in case of problems in the upgrade process.

Starting a TFTP Server

The switch ships with Intel Device View (see *Using Intel® Device View* for information about installing and using Intel Device View). To activate the TFTP server, choose Tools and then choose TFTP Server.

Once the TFTP server is running, click the Server Dir. button. Verify that the active directory is Program Files\Intel\Intel Device View\Firmware. Make sure that both the BootROM image (a file named **ngboot*nn*.bin**) and the firmware image (a file named **v*nnnn*b*nn*.tfp**) are in this directory.

Upgrading the BootROM

The BootROM image is always backward compatible with older versions of the firmware, so you can upgrade your BootROM before you upgrade your firmware without losing switch functionality.

To upgrade the BootROM image:

1. Connect an ethernet cable between the switch and a workstation that is on the same network or subnet as the switch. Use terminal emulation software, such as HyperTerminal, to connect to the switch and log in. By default, a login of **admin** with no password is provided
2. Save your existing configuration to disk using this command. Choose a filename you will remember easily. TFTPserverIP is the IP address of your TFTP server. You can find this in the lower left-hand corner of the TFTP server window.

```
upload configuration <TFTPserverIP> <filename>
```

3. Reset the switch to factory defaults using this command:

```
unconfig switch all
```

4. Log into switch and set the IP address of the switch to a valid IP address on your network.

```
configure vlan default ipaddress <ip address>
<mask>
```

5. Save this configuration to the primary database.

```
save configuration primary
```

6. Download BootROM 6.5 to the switch from your TFTP server.

```
download bootrom <TFTPserverIP> ngboot<nn>.bin
```

7. Reboot the switch and log back on.

```
reboot
```

Upgrading the Firmware

To upgrade the firmware on the switch:

1. Download the latest image from your TFTP server.

```
download image <TFTPserverIP> v<nnn>b<nn>.tfp
primary
```

2. Verify that primary image is now at the latest version and that the secondary image is still at the older version:

```
show switch
```

3. Save this configuration in the primary configuration database:

```
save configuration primary
```

4. Then reboot the switch, and log back into the switch.

```
reboot
```

5. Verify that the switch is now using the latest version of the BootROM and firmware:

```
show version
```

6. Download your saved configuration back onto the switch. <File-name> is the name of the configuration file you saved earlier before downloading the new BootROM.

```
download configuration <TFTPserverIP> <filename>
primary
```

Downgrading Your Switch

Assuming you have followed the upgrade instructions correctly, these steps return to your previous firmware and configuration files:

- Activate the previous image in the secondary image space using the command:

  ```
  use image secondary
  ```

- To configure the switch to access the secondary configuration (assuming you have set up the older version as the secondary configuration) use the command:

  ```
  use config secondary
  ```

- Verify that the above procedures were completed successfully with the command:

```
show switch
```

- Reboot the switch using the **`reboot`** command.

If you have followed upgrade instructions, your original configuration should be operational.

If you did not have an older configuration, you may perform a minimal configuration for the switch through the command line interface (CLI) sufficient to TFTP download the configuration file generated during the upgrade procedure.

Using TFTP to Upload the Configuration

You can upload the current configuration to a TFTP server on your network. The uploaded ASCII file retains the command-line interface (CLI) format. This allows you to:

- Modify the configuration using a text editor, and later download a copy of the file to the same switch, or to one or more different switches.
- Send a copy of the configuration file to Intel Customer Support for problem-solving.
- Automatically upload the configuration file every day, so that the TFTP server can archive the configuration daily. Because the filename is not changed, the configured file stored in the TFTP server is overwritten every day.

To upload the configuration, use the command:

```
upload configuration [<ipaddress> | <hostname>]
<filename> {every <time>}
```

where:

- **`ipaddress`** is the IP address of the TFTP server.
- **`hostname`** is the hostname of the TFTP server. (You must enable DNS to use this option.)
- **`filename`** is the name of the ASCII file. The filename can be up to 255 characters long, and cannot include any spaces, commas, quotation marks, or special characters.
- **`every <time>`** specifies the time of day you want the configuration automatically uploaded on a daily basis. If not

specified, the current configuration is immediately uploaded to the TFTP server.

To cancel a scheduled configuration upload, use the command:

```
upload configuration cancel
```

Using TFTP to Download the Configuration

To modify the switch configuration, you can download ASCII files that contain CLI commands to the switch. Three types of configuration scenarios can be downloaded:

- Complete configuration
- Incremental configuration
- Scheduled incremental configuration

You can find a TFTP Server Utility in Intel Device View under the Tools menu.

Downloading a Complete Configuration

Downloading a complete configuration replicates or restores the entire configuration to the switch. You typically use this type of download with the **upload config** command, which generates a complete switch configuration in an ASCII format. As part of the complete configuration download, the switch is automatically rebooted.

To download a complete configuration, use this command:

```
download configuration [<hostname | ip_address>]
<filename>
```

After you download the ASCII configuration using TFTP, you are prompted to reboot the switch. The downloaded configuration file is stored in current switch memory during the rebooting process, and is not retained if the switch has a power failure.

When the switch completes booting, it treats the downloaded configuration file as a script of CLI commands, and automatically executes the commands. If your CLI connection is through a Telnet connection (and not the console port), your connection is terminated when the switch reboots, but the command executes normally.

Downloading an Incremental Configuration

You can make a partial or incremental change to the switch configuration using downloaded ASCII files that contain CLI commands. The switch interprets these commands as a script of CLI commands. They take effect at the time of the download, without requiring a reboot of the switch.

To download an incremental configuration, use this command:

```
download configuration <hostname | ip_address>
<filename> {incremental}
```

Scheduled Incremental Configuration Download

You can schedule the switch to download a partial or incremental configuration on a regular basis. You can use this feature to update the switch configuration regularly from a centrally administered TFTP server. As part of the scheduled incremental download, you can optionally configure a backup TFTP server.

To configure the primary and/or secondary TFTP server and filename, use this command:

```
configure download server [primary | secondary]
<hostname | ip_address> <filename>
```

To enable scheduled incremental downloads, use this command:

```
download configuration every <hour> <min>
```

To display scheduled download information, use this command:

```
show switch
```

To cancel scheduled incremental downloads, use this command:

```
download configuration cancel
```

Remember to Save

Regardless of the download option used, configurations are downloaded into switch runtime memory only. The configuration is saved only when the **save** command is issued, or if the configuration file itself includes the **save** command.

If the configuration currently running in the switch does not match the configuration that the switch used when it originally booted, an asterisk (*) appears before the command line prompt when using the CLI.

Accessing BootROM

The BootROM of the switch initializes certain important switch variables during the boot process. In the event the switch does not boot properly, you can access some boot option functions through a special BootROM menu.

Interaction with the BootROM menu is only required under special circumstances, and should be done only under the direction of Intel Customer Support.

To access the BootROM menu, follow these steps:

- Attach a serial cable to the console port of the switch, as described in Chapter 3, Accessing the Switch.
- Attach the other end of the serial cable to a terminal or terminal emulator.
- Power cycle the switch while pressing the spacebar on the keyboard of the terminal.
- When the **BootROM->** prompt appears, release the spacebar. You can open a Help menu by pressing **h**.

  Options in the menu include:

  - Selecting the image to boot from
  - Booting to factory default configuration
  - Performing a serial download of an image

For example, to change the image that the switch boots from in flash memory:

- Press `1` for the image stored in primary, or

- Press **2** for the image stored in secondary.

Then, press the **f** key to boot from newly selected on-board flash memory.

To boot to factory default configuration:

- Press the **d** key for default, and
- Press the **f** key to boot from the configured on-board flash.

Boot Option Commands

Table 20.1 lists the CLI commands associated with switch boot options. For further command options, press the Tab key in the command line interface.

Table 20.1: Boot Option Commands

Command	Description
configure download server [primary | secondary] <hostname | ipaddress> <filename>	Configures the TFTP server(s) used by a scheduled incremental configuration download.
download bootrom [<ipaddress> | <hostname>] <filename>	Downloads a BootROM image from a TFTP server. The downloaded image replaces the BootROM in the onboard flash memory. ***Caution*** *If this command does not complete successfully, it could prevent the switch from booting.*
download configuration <hostname | ipaddress> <filename> {incremental}	Downloads a complete configuration. Use the **incremental** keyword to specify an incremental configuration download.
download configuration cancel	Cancels a scheduled configuration download.
download configuration every <hour> <min>	Schedules a configuration download. Specify the hour using a 24-hour clock, where the range is 0 to 23.

Table 20.1: Boot Option Commands (continued)

Command	Description
download image [<ipaddress> | <hostname>] <filename> {primary | secondary}	Downloads a new image from a TFTP server over the network. If parameters are not specified, the image is saved to the current image.
reboot {time <date> <time> | cancel}	Reboots the switch on the date and time specified. If you do not specify a reboot time, the reboot happens immediately following the command, and any scheduled reboots are cancelled. To cancel a scheduled reboot, use the `cancel` option.
save {configuration} {primary | secondary}	Saves the current configuration to nonvolatile (more permanent) storage. You can specify the primary or secondary configuration area. If not specified, the configuration is saved to the primary configuration area.
show configuration	Displays the current configuration to the terminal. You can then capture the output and store it as a file.
upload configuration [<ipaddress> | <hostname>] <filename> {every <time>}	Uploads the current run-time configuration to the specified TFTP server. If every `time` is specified, the switch automatically saves the configuration to the server once per day, at the specified time. If the time option is not specified, the current configuration is immediately uploaded.
upload configuration cancel	Cancels a scheduled configuration upload.

Table 20.1: Boot Option Commands (continued)

Command	Description
use configuration [primary | secondary]	Configures the switch to use a particular configuration on the next reboot. Options include the primary configuration area or the secondary configuration area.
use image [primary | secondary]	Configures the switch to use a particular image on the next reboot.

A Technical Specifications and Supported Limits

Technical Specifications

For IEEE standards information refer to http://standards.ieee.org

The following table lists specifications for the Intel® NetStructure™ 480T routing switch.

Table A.1: Specifications

Physical Dimensions	Height: 3.5 inches x Width: 17.36 inches x Depth: 19.20 inches Weight: with single PSU: 21.7 lbs with dual PSU: 27.4 lbs
Environmental Requirements	
Operating Temperature	0° to 40° C
Storage Temperature	-25 to 70° C
Operating Humidity	5% to 95% relative humidity, noncondensing
Standards	EN60068 (IEC68)
Certification Marks	

Table A.1: Specifications

CE	CE (European Community)
TÜV Rheinland GS	TUV/GS (German Notified Body)
✓	C-Tick (Australian Communication Authority)
c UL us LISTED	Underwriters Laboratories (USA and Canada)
Safety	
Agency Certifications	UL 1950 3rd Edition, listed cUL listed to CSA 22.2#950 TUV GS mark safety approval to the following EN standards: EN60950:1992/A3:+A1 +A2 +A3 +A4 +A11 EN60825-1; 1994, + All
Electromagnetic Compatibility	FCC part 15 Class Ices003 Issue3 Class A VCCI Class A EN55022:1998 Class A EN55024:1998 C-Tick mark to AS/NZS 3548:1997 Class A RRL (Korea) BSMI (Taiwan) CNS13438: 1997
Heat Dissipation	265W maximum (904.82 BTU/hr maximum)
Power Supply	
AC Line Frequency	47 Hz to 63 Hz
Input Voltage Options	90 VAC to 264 VAC, auto-ranging
Current Rating	100-120/200-240 VAC 4.0/2.0 A

Supported Standards, RFCs and Protocols

Table A.2: Supported Standards, RFCs and Protocols

RFCs, Standards, and Protocols
RFC 1058 RIP
RFC 1723 RIP v2
RFC 1112 IGMP
RFC 2236 IGMP v2
DVMRP v3 - Draft IETF DVMRP v3-07
PIM-DM v2 - Draft IETF PIM-DM v2-dm-01
RFC 2362 PIM-SM
RFC 1587-NSSA option
RFC 2178 OSPF
RFC 1122 Host requirements
IEEE 802.1D-1998 (802.1p) Packet priority
IEEE 802.1Q VLAN tagging
IEEE 802.3u 100 Mbps Ethernet
IEEE 802.3z 1 Gbps Ethernet
IEEE 802.3ab 1 Gbps Ethernet on Cat 5 UTP
IEEE 802.3ac Frame Extension for VLAN tagging on Ethernet
IEEE 802.3ad Link Aggregation
IEEE 802.3x Full-Duplex Operation/Flow Control on Ethernet
IEEE 802.1d Spanning Tree Protocol
RFC 1965 - Autonomous System Confederations for BGP
RFC 1771Border Gateway Protocol (BGP-4)
RFC 1966 - BGP Route Reflection
RFC 1997 - BGP Communities Attribute
RFC 1745 - BGP/OSPF
RFC 2113 - IP Router Alert Option
RFC 1256 Router discovery protocol
RFC 1812 IP router requirement
RFC 783 TFTP
RFC 1542 BootP
RFC 854 Telnet
RFC 768 UDP
RFC 791 IP
RFC 792 ICMP
RFC 793 TCP
RFC 826 ARP
RFC 2068 HTTP
RFC 2131 BootP/DHCP relay
RFC 2030 - Simple Network Time Protocol
IPX§ RIP/SAP Router specification
IPX SNAP
IP RIP v1 and v2
IP Multinetting
NetBEUI
AppleTalk§ protocol
Enterprise Standby Router Protocol (ESRP)
Management and Security
RFC 1157 SNMP v1/v2c
RFC 1213 MIB II
RFC 1354 IP forwarding table MIB
RFC 1493 Bridge MIB
RFC 2037 Entity MIB
RFC 1573 Evolution of Interface
RFC 1643 Ethernet MIB
RFC 1757 Four groups of RMON
RFC 2021 RMON probe configuration
RFC 2239 802.3 MAU MIB
RFC 1724 RIP v2 MIB
Enterprise MIB
HTML and Telnet management
RFC 2138 RADIUS

Supported Limits

The table below summarizes tested metrics for various features on the 480T routing switch. These metrics are laboratory results and are for reference and comparison only.

Table A.3: Supported Limits

Metric	Description	Limit
Access Profiles	Used by SNMP, Telnet, Vista Web interface, and Routing Access Policies.	128
Access Profile entries	Used by SNMP, Telnet, Vista Web interface, and Routing Access Policies.	256
Access List rules	Maximum number of Access Lists in which all rules utilize all available options.	worst case: 255
Telnet - number of sessions	Maximum number of simultaneous Telnet sessions.	8
SNMP - Trap receivers	Maximum number of SNMP trap receiver stations supported.	16
Syslog servers	Maximum number of simultaneous syslog servers that are supported.	4
Jumbo Frame size	Maximum size supported for Jumbo frames, including the CRC.	9216
VLANs	Includes all VLANs, sub-VLANs, super-VLANs	3000
IP Router interfaces	Maximum number of VLANs performing IP routing; excludes sub-VLANs.	512

Table A.3: Supported Limits

MAC-based VLANs – MAC addresses	Maximum number of MAC addresses that can be downloaded to the switch when using MAC-based VLANs.	7000
Protocol-sensitive VLANs – active protocol filters	The number of simultaneously active protocol filters in the switch.	15
Spanning Tree - Max STPDs	Maximum number of Spanning Tree Domains.	64
Spanning Tree – Maximum number of ports	Maximum number of ports that can participate in a single Spanning Tree Domain.	Same as available physical ports.
IP Static Routes	Maximum number of permanent IP routes.	1024
IP route sharing entries	Maximum number of IP routes used in route sharing calculations. This includes static routes and OSPF ECMP.	8
IP Static ARP entries	Maximum number of permanent IP static ARP entries supported.	512
Static IP ARP Proxy entries	Maximum number of permanent IP ARP proxy entries.	512
Static MAC FDB entries	Maximum number of permanent MAC entries configured into the FDB.	256
UDP profiles	Number of profiles that can be created for UDP forwarding.	10
UDP profile entries	Number of entries within a single UDP profile.	16
ESRP Route-track entries	Maximum number of routes that can be tracked by ESRP.	256

Table A.3: Supported Limits

ESRP – number of instances	Maximum number of ESRP-supported VLANs for a single switch.	64
ESRP – number of ESRP groups	Maximum number of ESRP groups within a broadcast domain.	4
ESRP – number of VLANs in a single ESRP domain	Maximum number of VLANs that can be joined to a single ESRP instance through an ESRP domain.	256 default; 300 max
FDB – Maximum number of L2/L3 entries	Maximum number of MAC addresses.	128,000
Mirroring – Mirrored ports	Maximum number of ports that can be mirrored to the mirror port.	8
Mirroring – number of VLANs	Maximum number of VLANs that can be mirrored to the mirror port.	8
RIP-learned routes	Maximum number of RIP routes supported without aggregation.	8000
RIP interfaces on a single router	Recommended maximum number of RIP-routed interfaces on a switch.	384
OSPF areas	As an ABR, how many OSPF areas are supported within the same switch.	8
OSPF routes	Recommended maximum number of routes contained in an OSPF LSDB.	30,000
OSPF routers in a single area	Recommended maximum number of routers in a single OSPF area.	40
OSPF interfaces on a single router	Recommended maximum number of OSPF-routed interfaces on a switch.	384

Table A.3: Supported Limits

OSPF virtual links	Maximum number of OSPF virtual links supported.	32
BGP routes	Maximum number of routes contained in the BGP route table.	500,000
BGP peers	Maximum number of BGP peers on a single router.	64
Policy-Based Routing	Maximum number of policy-based routes that can be stored on a switch.	64
WCR - Max number of redirection rules	Maximum number of rules that can point to the same or separate groups of Web cache servers.	64 (8 servers)
SLB - Max number of simultaneous sessions	For Transparent and Translational and GoGo modes respectively.	500,000/ 500,000/ unlimited
SLB - Max number of VIPs	For Transparent and Translational and GoGo modes respectively.	1000/1000/ unlimited
SLB - Max number of Pools	For Transparent and Translational (does not apply to GoGo mode).	256/256
SLB - Max number of Nodes per Pool	For Transparent and Translational (does not apply to GoGo mode).	256/256
SLB - Max number of physical servers per group	Applies to GoGo mode only; a group shares any number of common VIPs.	8
IPX[§] static routes and services (RIP and SAP)	Maximum number of static IPX RIP route and IPX SAP entries.	64 for each
IPX dynamic routes and services	Maximum recommended number of dynamically learned IPX RIP routes and SAP entries.	2000 for each

Table A.3: Supported Limits

IPX Router interfaces	Maximum number of IPX router interfaces.	256
IPX Access control lists	Maximum number of access lists in which all rules utilize all available options.	worst case: 255

B Troubleshooting

If you encounter problems when using the Intel® NetStructure™ 480T routing switch, this appendix may be helpful. If you have a problem not listed here or in the “Late Breaking News,” contact your local technical support representative (see "Intel Customer Support" on page 491).

LEDs

Why doesn’t the power LED light?

- Check that the power cable is firmly connected to the device and to the supply outlet.

Why does the MGMT LED light *orange* when powering on?

- The device has failed its Power-On Self Test (POST). Contact your supplier for advice.

A link is connected, but the Status LED does not light. Why?

- Check that all connections are secure.
- Make sure cables are free from damage.
- Make sure the devices at both ends of the link are powered-up.
- Ensure both ends of the 1000 Mbps link are set to the same autonegotiation state.
- Both sides of the 1000 Mbps link must be enabled or disabled. If the two are different, typically the side with autonegotiation disabled will have

the link LED lit, and the side with autonegotiation enabled will not have the LED lit.

- The default configuration for a 1000 Mbps port is autonegotiation enabled. Verify by using this command:

```
show port config
```

Why won't the switch power on?

- The 480T routing switch uses a digital power supply with surge protection. During a power surge, the protection circuits turn off the power supply.
- To reset, unplug the switch for 1 minute, plug it back in, and restart the switch.
- If this does not work, try using a different power source (different power strip or outlet) and power cord.

Using the Command-Line Interface

Why won't the initial *Welcome* prompt display?

- Check that your terminal or terminal emulator is correctly configured.
- For console port access, you may need to press Enter several times before the welcome prompt appears.
- Check the settings on your terminal or terminal emulator. The settings are 9600 baud, 8 data bits, 1 stop bit, no parity, XON/OFF flow control enabled.

Why won't the SNMP Network Manager access the device?

- Check that the device IP address, subnet mask, and default gateway are correctly configured, and that the device has been reset.
- Check that the device IP address is correctly recorded by the SNMP Network Manager (refer to your user documentation for the Network Manager).
- Check that the community strings configured for the system and Network Manager are the same.
- Check that SNMP access was not disabled for the system.

Why won’t the Telnet workstation access the device?

- Check that the device IP address, subnet mask and default gateway are configured correctly, and that the device has been reset.
- Ensure that you enter the IP address of the switch correctly when invoking the Telnet facility.
- Check that Telnet access was not disabled for the switch. If you attempt to log in and the maximum number of Telnet sessions are being used, you should receive an error message indicating so.

Why is it that traps are not received by the SNMP Network Manager?

- Check that the SNMP Network Manager's IP address and community string are configured correctly
- Check that the IP address of the Trap Receiver is configured properly on the system.

The SNMP Network Manager or Telnet workstation can no longer access the device. Why?

- Check that Telnet access or SNMP access is enabled.
- Check that the port through which you are trying to access the device has not been disabled. If it is enabled, check the connections and network cabling at the port.
- Check that the port through which you are trying to access the device is in a correctly configured VLAN.
- Try accessing the device through a different port. If you succeed, a problem with the original port is indicated. Examine the connections and cabling.
- A network problem may be blocking your access to the device over the network. Try accessing the device through the console port.
- Check that the community strings configured for the device and the Network Manager are the same.
- Check that SNMP access was not disabled for the system.

Why do permanent entries remain in the Forwarding Database (FDB)?

- If you have made a permanent entry in the FDB (which requires you to specify the VLAN where it belongs and then delete the VLAN), the FDB entry will remain. Though harmless, you must manually delete the entry from the FDB if you want to remove it.

How do I remove unused *default* and *static* routes?

- If you have defined static or default routes, those routes will remain in the configuration, independent of whether the VLAN and VLAN IP address that used them remains. You should manually delete the routes if no VLAN IP address is capable of using them.

What if I forget my password and cannot log in?

- If you are not an administrator, another user having administrator access level can log in, delete your user name, and create a new user name for you, with a new password.
- Alternatively, another user having administrator access level can log in and reset the device to factory defaults. This will return all configuration information (including passwords) to the initial values.
- If no one knows a password for an administrator level user, contact Intel Customer Support. See “Intel Customer Support” on page 461.

Port Configuration

What if no link light shows on a 100/1000 Base port?

- If patching from a hub or switch to another hub or switch, ensure that you are using a CAT5 (Category 5) cable. A crossover cable is recommended and is required for some configurations.

What if I’m receiving excessive RX CRC errors?

- When a device with autonegotiation disabled is connected to a 480T routing switch that has autonegotiation enabled, the switch links at the correct speed, but in half-duplex mode.

  The switch 100/1000 Mbps physical interface uses a method called parallel detection to access the link. Because the other network device is not participating in autonegotiation (and does not advertise its capabilities), parallel detection on the switch is only able to sense 100 Mbps versus 1000 Mbps speed, and not the duplex mode. Therefore, the switch establishes the link in half-duplex mode using the correct speed.

- To establish a full-duplex link either force it at both sides, or run autonegotiation on both sides (using full-duplex as an advertised capability, the default setting).

Always verify that the switch and the network device match in configuration for speed and duplex.

- A mismatch of duplex mode between the 480T routing switch and another network device will cause poor network performance. View statistics using:

  ```
  show port rx
  ```

  If it displays a constant increment of CRC errors, it is a duplex mismatch between devices, rather than a problem with the 480T routing switch.

What if no link light shows on a 1000 Mbps fiber port?

- Check to ensure that the transmit fiber goes to the receive fiber side of the other device, and vice-versa.
- The switch has autonegotiation set to on by default for 1000 Mbps ports. Set these ports to auto off (using the command `configure port <port #> auto off`) if you are connecting to devices that do not support autonegotiation.
- Ensure that you are using multi-mode fiber (MMF) when using a 1000BASE-SX GBIC, and single mode fiber (SMF) when using a 1000BASE-LX GBIC. 1000BASE-SX does not work with SMF. 1000BASE-LX works with MMF, but requires you to use a mode conditioning patchcord (MCP).

OSPF (Open Shortest Path First)

When setting up OSPF areas, it indicates the area must be in an IP-type format. That differs from some non-Intel equipment. How do I convert an OSPF area into an IP-type format?

The 480T routing switch must have the OSPF area ID input in IP dotted decimal notation. Some non-Intel equipment may show this as a whole number. To convert OSPF whole numbers to dotted decimal notation:

- Convert the non-IP type format using a decimal to binary converting method, for example, to convert 400 decimal into binary (110010000). The binary number needs to show 32 digits, representing the digits of the 4 octets in the IP-type format. 110010000 binary = 00000000.00000000.0000001.10010000 as broken into octets.

- Then convert each octet into a decimal value. (for example, 00000000.00000000.0000001.10010000 = 0.0.1.144).
- Therefore, 400 = 0.0.1.144

VLANs

What if I can't add a port to a VLAN?

- If you attempt to add a port to a VLAN and get an error message similar to:

  ```
  localhost:7 # configure vlan marketing add port
  1,2
  ```

  ```
  ERROR: Protocol conflict on port 1
  ```

  you already have a VLAN using untagged traffic on a port. You can only configure one VLAN using untagged traffic on a single physical port.

- Verify VLAN configuration by using this command:

  ```
      show vlan <name>
  ```

- The solution for this error is to remove ports 1 and 2 from the VLAN currently using untagged traffic on those ports. If this were the default VLAN, you would use the command:

  ```
  localhost:23 # configure vlan default del port
  1,2
  ```

  This allows you to enter the previous command (without getting an error message) as:

  ```
  localhost:26 # configure vlan red add port 1,2
  ```

VLAN Names

There are restrictions on VLAN names. They cannot contain white spaces and cannot start with a numeric value unless you use quotation marks around the name.

If a name contains white spaces, starts with a number, or contains non-alphabetical characters, you must use quotation marks whenever referring to the VLAN name.

What if 802.1Q links do not work correctly?

- VLAN names are only locally significant through the command-line interface. For two switches to communicate across an 802.1Q link, the VLAN ID for the VLAN on one switch should have a corresponding VLAN ID for the VLAN on the other switch.
- If you are connecting to a third-party device and have checked that the VLAN IDs are the same, the Ethertype field used to identify packets as 802.1Q packets may differ between the devices. The default value used by the switch is 8100. If the third-party device differs from this and cannot be changed, you can change the 802.1Q Ethertype used by the switch with this command:

  ```
  configure dot1p ethertype <ethertype>
  ```

  Changing this parameter changes how the system recognizes all tagged frames it receives, as well as the value it inserts in all tagged frames it transmits.

VLANs, IP Addresses and Default Routes

The system can have an IP address for each configured VLAN. It is necessary to have an IP address associated with a VLAN if you intend to manage (Telnet, SNMP, ping) through that VLAN or route IP traffic.

You can also configure multiple default routes for the system. The system first tries the default route with the lowest cost metric.

STP

I have connected an endstation directly to the switch, but the endstation fails to boot correctly. Why?

- The switch has STP enabled, and the endstation is booting before the STP initialization process is complete.
- Verify that STP has been disabled for that VLAN, or turn off STP for the switch ports of the endstation (and devices to which it is attempting to connect). Then reboot the endstation.

Why does the switch keep aging out endstation entries in the switch Forwarding Database (FDB)?

- Reduce the number of topology changes by disabling STP on those systems that do not use redundant paths.
- Specify that the endstation entries are static or permanent.

ESRP

Why am I having trouble with interoperability between switches using different versions of firmware running ESRP?

- We recommend that all switches running ESRP use the same version of firmware.
- If mixing of firmware versions becomes necessary on the network, problems may arise using some of the newer ESRP features. Contact customer support for details.

Troubleshooting Tools

Debug Tracing

The debug commands should only be used when advised by Intel technical personnel.

The local management software includes a debug-tracing facility for the switch. The command can be applied to one or all VLANs:

```
show debug-tracing {vlan <name>}
```

TOP Command

The `top` command activates a utility that indicates microprocessor utilization by process.

C Regulatory Information

Compliance statements

Each of the following compliance statements applies only to products that bear the mark or text required by the appropriate certification agency.

FCC Part 15 Compliance Statement

This product has been tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment.

This product generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference; in which case, the user will be required to correct the interference at his own expense.

NOTE: This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) This device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

CAUTION If you make any modification to the equipment not expressly approved by Intel, you could void your authority to operate the equipment.

Canada Compliance Statement (Industry Canada)

Cet appareil numérique respecte les limites bruits radioélectriques applicables aux appareils numériques de Classe A prescrites dans la norme sur le matériel brouilleur: "Appareils Numériques", NMB-003 édictée par le Ministre Canadien des Communications.

This digital apparatus does not exceed the Class A limits for radio noise emissions from digital apparatus set out in the interference-causing equipment standard entitled: "Digital Apparatus," ICES-003 of the Canadian Department of Communications.

CE Compliance Statement

This certifies that the Intel® NetStructure™ 480T routing switch complies with the EU Directive, 89/336/EEC, using the EMC standards EN55022 (Class A) and EN50082-1. This product also complies with the EU Directive, 73/23/EEC, using the safety standard EN60950. In addition, this product complies with the EU Standards EN61000-3-2 and EN61000-3-3.

CISPR 22 Statement

WARNING This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

Taiwan Class A EMI Statement

警告使用者：

這是甲類的資訊產品， 在居住的環境中使用時， 可能會造成射頻干擾， 在這種情況下， 使用者會被要求採取某些適當的對策。

VCCI Statement

Class A ITE

この装置は、クラス A 情報技術装置です。この装置を家庭環境で使用すると電波障害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるように要求されることがあります。 VCCI-A

This is a Class A product based on the standard of the Voluntary Control Council for Interference by Information Technology Equipment (VCCI). If this equipment is used in a domestic environment, radio disturbance may arise. When such trouble occurs, the user may be required to take corrective actions.

Warnings

WARNING

This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

Internal access to the Intel NetStructure 480T routing switch is intended only for qualified service personnel. Do not remove any covers. There are no user serviceable parts inside.

WARNING

Choose a site that is:

- Clean and free of airborne particles (other than normal room dust).
- Well ventilated and away from sources of heat including direct sunlight.
- Away from sources of vibration or physical shock.
- Isolated from strong electromagnetic fields produced by electrical devices.
- In regions that are susceptible to electrical storms, we recommend you plug your system into a surge suppressor and disconnect telecommunication lines to your modem during an electrical storm.
- Provided with a properly grounded wall outlet.

Do not attempt to modify or use the supplied AC power cord if it is not the exact type required.

AVERTISSEMENT

L’emplacement choisi doit être:

- Propre et dépourvu de poussière en suspension (sauf la poussière normale).
- Bien aéré et loin des sources de chaleur, y compris du soleil direct.
- A l’abri des chocs et des sources de ibrations.
- Isolé de forts champs magnétiques géénérés par des appareils électriques.
- Dans les régions sujettes aux orages magnétiques il est recomandé de brancher votre système à un supresseur de surtension, et de débrancher toutes les lignes de télécommunications de votre modem durant un orage.
- Muni d’une prise murale correctement mise à la terre.

Ne pas utiliser ni modifier le câble d’alimentation C. A. fourni, s’il ne correspond pas exactement au type requis.

WARNUNG

Der entwickelt. Der Standort sollte:

- sauber und staubfrei sein (Hausstaub ausgenommen);
- gut gelüftet und keinen Heizquellen ausgesetzt sein (einschließlich direkter Sonneneinstrahlung);
- keinen Erschütterungen ausgesetzt sein;
- keine starken, von elektrischen Geräten erzeugten elektromagnetischen Felder aufweisen;

- in Regionen, in denen elektrische Stürme auftreten, mit einem Überspannungsschutzgerät verbunden sein; während eines elektrischen Sturms sollte keine Verbindung der Telekommunikationsleitungen mit dem Modem bestehen;
- mit einer geerdeten Wechselstromsteckdose ausgerüstet sein.

Versuchen Sie nicht, das mitgelieferte Netzkabel zu ändern oder zu verwenden, wenn es sich nicht um genau den erforderlichen Typ handelt.

AVVERTENZA

Scegliere una postazione che sia:

- Pulita e libera da particelle in sospensione (a parte la normale polvere presente nell'ambiente).
- Ben ventilata e lontana da fonti di calore, compresa la luce solare diretta.
- Al riparo da urti e lontana da fonti divibrazione.
- Isolata dai forti campi magnetici prodotti da dispositivi elettrici.
- In aree soggette a temporali, è consigliabile collegare il sistema ad un limitatore di corrente. In caso di temporali, scollegare le linee di comunicazione dal modem.
- Dotata di una presa a muro correttamente installata.

Non modificare o utilizzare il cavo di alimentazione in c. a. fornito dal produttore, se non corrisponde esattamente al tipo richiesto.

ADVERTENCIAS

Escoja un lugar:

- Limpio y libre de partículas en suspensión (salvo el polvo normal)
- Bien ventilado y alejado de fuentes de calor, incluida la luz solar directa.
- Alejado de fuentes de vibración.
- Aislado de campos electromagnéticos fuertes producidos por dispositivos eléctricos.
- En regiones con frecuentes tormentas eléctricas, se recomienda conectar su sistema a un eliminador de sobrevoltage y desconectar el módem de las líneas de telecomunicación durante las tormentas.
- Previsto de una toma de tierra correctamente instalada.

No intente modificar ni usar el cable de alimentación de corriente alterna, si no se corresponde exactamente con el tipo requerido.

Limited Hardware Warranty

Intel warrants to the original owner that the hardware product delivered in this package will be free from defects in material and workmanship for three (3) years following the latter of: (i) the date of purchase only if you register by returning the registration card as indicated thereon with proof of purchase; or (ii) the date of manufacture; or (iii) the registration date if by electronic means provided such registration occurs within thirty (30) days from purchase. This warranty does not cover the product if it is damaged in the process of being installed. Intel recommends that you have the company from whom you purchased this product install the product.

INTEL RESERVES THE RIGHT TO FILL YOUR ORDER WITH A PRODUCT CONTAINING NEW OR REMANUFACTURED COMPONENTS. THE ABOVE

WARRANTY IS IN LIEU OF ANY OTHER WARRANTY, WHETHER EXPRESS, IMPLIED OR STATUTORY, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTY OF NONINFRINGEMENT OF INTELLECTUAL PROPERTY, MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR ANY WARRANTY ARISING OUT OF ANY PROPOSAL, SPECIFICATION, SAMPLE OR OTHERWISE.

This warranty does not cover replacement of products damaged by abuse, accident, misuse, neglect, alteration, repair, disaster, improper installation or improper testing. If the product is found to be otherwise defective, Intel, at its option, will replace or repair the product at no charge except as set forth below, provided that you deliver the product along with a return material authorization (RMA) number either to the company from whom you purchased it or to Intel (North America only). If you ship the product, you must assume the risk of damage or loss in transit. You must use the original container (or the equivalent) and pay the shipping charge. Intel may replace or repair the product with either new or remanufactured product or parts, and the returned product becomes Intel's property. Intel warrants the repaired or replaced product to be free from defects in material and workmanship for a period of the greater of: (i) ninety (90) days from the return shipping date; or (ii) the period of time remaining on the original three (3) year warranty.

This warranty gives you specific legal rights and you may have other rights which vary from state to state. All parts or components contained in this product are covered by Intel's limited warranty for this product; the product may contain fully tested, recycled parts, warranted as if new. For warranty information call one of the numbers below.

Returning a Defective Product (RMA)

Before returning any product, contact an Intel Customer Support Group and obtain an RMA number by calling:

North America only: (916) 377-7000
Other locations: Return the product to the place of purchase.

If the Customer Support Group verifies that the product is defective, they will have the Return Material Authorization Department issue you an RMA number to place on the outer package of the product. Intel cannot accept any product without an RMA number on the package.

LIMITATION OF LIABILITY AND REMEDIES

INTEL SHALL HAVE NO LIABILITY FOR ANY INDIRECT OR SPECULATIVE DAMAGES (INCLUDING, WITHOUT LIMITING THE FOREGOING, CONSEQUENTIAL, INCIDENTAL AND SPECIAL DAMAGES) ARISING FROM THE USE OF OR INABILITY TO USE THIS PRODUCT, WHETHER ARISING OUT OF CONTRACT, NEGLIGENCE, TORT, OR UNDER ANY WARRANTY, IRRESPECTIVE OF WHETHER INTEL HAS ADVANCE NOTICE OF THE POSSIBILITY OF ANY SUCH DAMAGES, INCLUDING, BUT NOT LIMITED TO LOSS OF USE, INFRINGEMENT OF INTELLECTUAL PROPERTY, BUSINESS INTERRUPTIONS, AND LOSS OF PROFITS, NOTWITHSTANDING THE FOREGOING, INTEL'S TOTAL LIABILITY FOR ALL CLAIMS UNDER THIS AGREEMENT SHALL NOT EXCEED THE PRICE PAID FOR THE PRODUCT. THESE LIMITATIONS ON POTENTIAL LIABILITIES WERE AN ESSENTIAL ELEMENT IN SETTING THE PRODUCT PRICE. INTEL NEITHER ASSUMES NOR AUTHORIZES ANYONE TO ASSUME FOR IT ANY OTHER LIABILITIES.

Some states do not allow the exclusion or limitation of incidental or consequential damages, so the above limitations or exclusions may not apply to you.

Critical Control Applications: Intel specifically disclaims liability for use of the hardware product in critical control applications (including, for example only, safety or health care control systems, nuclear energy control systems, or air or ground traffic control systems) by Licensee or Sublicensees, and such use is entirely at the user's risk. Licensee agrees to defend, indemnify, and hold Intel harmless from and against any and all claims arising out of use of the hardware product in such applications by Licensee or Sublicensees.

Software: Software provided with the hardware product is not covered under the hardware warranty described above. See the applicable software license agreement which shipped with the hardware product for details on any software warranty.

Limited Hardware Warranty (Europe only)

Intel warrants to the original owner that the hardware product delivered in this package will be free from defects in material and workmanship for three (3) years following the latter of: (i) the date of purchase only if you register by returning the registration card as indicated thereon with proof of purchase; or (ii) the date of manufacture; or (iii) the registration date if by electronic means provided such registration occurs within thirty (30) days from purchase. This warranty does not cover the product if it is damaged in the process of being installed. Intel recommends that you have the company from whom you purchased this product install the product.

INTEL RESERVES THE RIGHT TO FILL YOUR ORDER WITH A PRODUCT CONTAINING NEW OR REMANUFACTURED COMPONENTS. THE ABOVE WARRANTY IS IN LIEU OF ANY OTHER WARRANTY, WHETHER EXPRESS, IMPLIED OR STATUTORY, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTY OF NONINFRINGEMENT OF INTELLECTUAL PROPERTY, SATISFACTORY QUALITY, FITNESS FOR A PARTICULAR PURPOSE, OR ANY WARRANTY ARISING OUT OF ANY PROPOSAL, SPECIFICATION, SAMPLE OR OTHERWISE.

This warranty does not cover replacement of products damaged by abuse, accident, misuse, neglect, alteration, repair, disaster, improper installation or improper testing. If the product is found to be otherwise defective, Intel, at its option, will replace or repair the product at no charge except as set forth below, provided that you deliver the product along with a return material authorization (RMA) number either to (a) the company from whom you purchased it or (b) to Intel, North America only (if purchased in Europe you must deliver the product to "(a)". If you ship the product, you must assume the risk of damage or loss in transit. You must use the original container (or the equivalent) and pay the shipping charge. Intel may replace or repair the product with either new or remanufactured product or parts, and the returned product becomes Intel's property. Intel warrants the repaired or replaced product to be free from defects in material and workmanship for a period of the greater of: (i) ninety (90) days from the return shipping date; or (ii) the period of time remaining on the original three (3) year warranty.

This warranty gives you specific legal rights and you may have other rights which vary from state to state. All parts or components contained in this product are covered by Intel's limited warranty for this product; the product may contain fully tested, recycled parts, warranted as if new. For warranty information call one of the numbers below.

Returning a Defective Product (RMA)

Before returning any product, contact an Intel Customer Support Group and obtain an RMA number by calling the non-toll free numbers below:

Country	Number	Language
France	+33 (0) 1 41 91 85 29	French

Country	Number	Language
Germany	+49 (0) 69 9509 6099	German
Italy	+39 (0) 2 696 33276	Italian
United Kingdom	+44 (0) 870 607 2439	English

If the Customer Support Group verifies that the product is defective, they will have the Return Material Authorization Department issue you an RMA number to place on the outer package of the product. Intel cannot accept any product without an RMA number on the package.

LIMITATION OF LIABILITY AND REMEDIES

INTEL SHALL HAVE NO LIABILITY FOR ANY INDIRECT OR SPECULATIVE DAMAGES (INCLUDING, WITHOUT LIMITING THE FOREGOING, CONSEQUENTIAL, INCIDENTAL AND SPECIAL DAMAGES) ARISING FROM THE USE OF OR INABILITY TO USE THIS PRODUCT, WHETHER ARISING OUT OF CONTRACT, NEGLIGENCE, TORT, OR UNDER ANY WARRANTY, IRRESPECTIVE OF WHETHER INTEL HAS ADVANCE NOTICE OF THE POSSIBILITY OF ANY SUCH DAMAGES, INCLUDING, BUT NOT LIMITED TO LOSS OF USE, INFRINGEMENT OF INTELLECTUAL PROPERTY, BUSINESS INTERRUPTIONS, AND LOSS OF PROFITS, NOTWITHSTANDING THE FOREGOING, INTEL'S TOTAL LIABILITY FOR ALL CLAIMS UNDER THIS AGREEMENT SHALL NOT EXCEED THE PRICE PAID FOR THE PRODUCT. THESE LIMITATIONS ON POTENTIAL LIABILITIES WERE AN ESSENTIAL ELEMENT IN SETTING THE PRODUCT PRICE. INTEL NEITHER ASSUMES NOR AUTHORIZES ANYONE TO ASSUME FOR IT ANY OTHER LIABILITIES.

Critical Control Applications: Intel specifically disclaims liability for use of the hardware product in critical control applications (including, for example only, safety or health care control systems, nuclear energy control systems, or air or ground traffic control systems) by Licensee or Sublicensees, and such use is entirely at the user's risk. Licensee agrees to defend, indemnify, and hold Intel harmless from and against any and all claims arising out of use of the hardware product in such applications by Licensee or Sublicensees.

Software: Software provided with the hardware product is not covered under the hardware warranty described above. See the applicable software license agreement which shipped with the hardware product for details on any software warranty.

This limited hardware warranty shall be governed by and construed in accordance with the laws of England and Wales. The courts of England shall have exclusive jurisdiction regarding any claim brought under this warranty.

Limitation de garantie du matériel (Europe)

Intel garantit au propriétaire original que le produit matériel livré dans le présent coffret est exempt de défaut matériel ou de fabrication pour une période de trois (3) ans à compter de la plus récente des dates suivantes : (i) la date d'achat uniquement si vous vous êtes inscrit en renvoyant la carte d'inscription de la façon indiquée, avec une preuve d'achat ; (ii) la date de fabrication ou (iii) la date d'inscription électronique à condition qu'elle ait lieu dans les 30 jours suivant l'achat. La présente garantie sera nulle si le produit matériel est endommagé lors de son installation. Intel recommande de faire installer le produit matériel par la société auprès de laquelle il a été acheté.

INTEL SE RESERVE LE DROIT DE VOUS LIVRER UN PRODUIT CONTENANT DES COMPOSANTS NOUVEAUX OU REPARES. CETTE GARANTIE REMPLACE TOUTES LES AUTRES GARANTIES, EXPRESSES, TACITES OU LEGALES, Y COMPRIS, MAIS SANS QUE CETTE ENUMERATION SOIT LIMITATIVE, LES GARANTIES CONCERNANT LE NON RESPECT DE LA PROPRIETE INTELLECTUELLE, LA QUALITE SATISFAISANTE, L'ADEQUATION POUR UN USAGE PARTICULIER, OU TOUTE AUTRE GARANTIE ISSUE DE TOUT AUTRE PROPOSITION, SPECIFICATION, ECHANTILLON OU AUTRE.

La présente garantie ne couvre pas le remplacement de produits matériels endommagés par abus, accident, mauvaise utilisation, négligence, altération, réparation, catastrophe, installation ou tests incorrects. Si le produit matériel s'avère défectueux pour une autre raison, Intel décidera de le remplacer ou de le réparer gratuitement, à l'exception des cas énumérés ci-après, à condition que le produit soit renvoyé avec un numéro d'autorisation de retour du matériel (ARM) à (a) la société auprès de laquelle il a été acheté ou (b) à Intel, en Amérique du Nord seulement (si l'achat a eu lieu en Europe vous devez le renvoyer à "(a)". Si vous expediéz le produit matériel, vous devez assumer le risque de dégâts ou de perte pendant le transport. Vous devez utiliser le coffret original (ou l'équivalent) et payer les frais de transport. Intel peut réparer le produit matériel ou le remplacer par un produit neuf ou remis à neuf, le produit renvoyé devenant la propriété d'Intel. Intel garantit que le produit matériel réparé ou de remplacement est exempt de défaut matériel ou de fabrication pendant la plus longue des périodes suivantes: (i) quatre-vingt-dix (90) jours à compter de la date de retour; ou (ii) la période encore couverte par la garantie originale de trois (3) ans.

La présente garantie vous accorde des droits juridiques spécifiques et vous pouvez également disposer d'autres droits variant d'un Etat à l'autre. Tous les composants ou pièces du produit matériel sont couverts par la garantie limitée d'Intel relative à ce dernier ; il peut contenir des pièces recyclées, entièrement testées et garanties comme neuves. Pour plus d'informations sur la garantie, appelez l'un des numéros énumérés ci-après.

Retour d'un produit défectueux (ARM)

Avant de retourner un produit matériel, contactez le service d'assistance à la clientèle Intel pour obtenir un numéro ARM.

Pays	Numéro	Langue
France	+33 (0) 1 41 91 85 29	Français
Allemagne	+49 (0) 69 9509 6099	Allemand
Italie	+39 (0) 2 696 33276	Italien
R.U.	+44 (0) 870 607 2439	Anglais

Si le service d'assistance confirme que le produit est défectueux, il demandera au Département d'autorisation de retour de matériel de vous attribuer un numéro ARM à indiquer sur l'emballage externe. Intel ne peut accepter aucun produit sans numéro ARM.

LIMITATION DE RESPONSABILITE ET DE RECOURS

INTEL DECLINE TOUTE RESPONSABILITE RELATIVE A DES DOMMAGES INDIRECTS OU SPECULATIFS (Y COMPRIS, SANS LIMITATION DES ELEMENTS CI-DESSUS, LES DOMMAGES CONSECUTIFS, ACCIDENTELS ET SPECIAUX) DECOULANT DE L'UTILISATION OU DE L'INCAPACITE D'UTILISER CE PRODUIT, DUS A UN CONTRAT, UNE NEGLIGENCE, UN TORT OU COUVERTS PAR TOUTE GARANTIE, MEME SI LA POSSIBILITE D'UN TEL DOMMAGE A DEJA ETE PORTEE A LA CONNAISSANCE D'INTEL, Y COMPRIS, MAIS SANS QUE CETTE ENUMERATION SOIT LIMITATIVE, UNE PRIVATION DE JOUISSANCE, UN NON RESPECT DE LA PROPRIETE INTELLECTUELLE, UNE INTERRUPTION DES ACTIVITES ET UN MANQUE A GAGNER . NONOBSTANT LA DECLARATION QUI PRECEDE, LA RESPONSABILITE GLOBALE DE INTEL

CONCERNANT TOUS LES LITIGES RELATIFS AU PRESENT ACCORD NE SERA PAS SUPERIEURE AU PRIX PAYE POUR LE PRODUIT. CES LIMITATIONS DE RESPONSABILITE POTENTIELLE ONT CONSTITUE UN FACTEUR DETERMINANT LORS DE LA FIXATION DU PRIX DU PRODUIT. INTEL N'ASSUME AUCUNE AUTRE RESPONSABILITE ET N'AUTORISE QUICONQUE A LE FAIRE EN SON NOM.

La garantie limitée du matériel est régie et interprétée par les lois en vigueur en Angleterre et au Pays de Galles. Les tribunaux anglais jouissent d'une juridiction exclusive en matière de litige concernant cette garantie.

Garanzia limitata sull'hardware (valida solo in Europa)

La garantie limitée du matériel est régie et interprétée par les lois en vigueur en Angleterre et au Pays de Galles. Les tribunaux anglais jouissent d'une juridiction exclusive en matière de litige concernant cette garantie.

Intel garantisce al proprietario originale che il prodotto hardware incluso in questo pacchetto è privo di difetti in materiale e in lavorazione per un periodo di tre (3) anni a partire dall'ultima data tra: (i) la data di acquisto, solo nel caso in cui l'utente effettua la registrazione tramite la scheda di registrazione, come indicato, accompagnata dalla prova di acquisto; oppure (ii) la data di fabbricazione; oppure (iii) la data di registrazione, se effettuata per via elettronica, a condizione che tale registrazione avvenga entro trenta (30) giorni dall'acquisto. Questa garanzia non copre il prodotto nel caso questo fosse danneggiato durante l'installazione. Intel raccomanda di fare installare il prodotto dall'azienda da cui il prodotto è stato acquistato.

INTEL SI RISERVA IL DIRITTO DI ONORARE L'ORDINAZIONE CON UN PRODOTTO CONTENENTE PARTI NUOVE O RIFABBRICATO. LA GARANZIA QUI SOPRA SOSTITUISCE QUALSIASI ALTRA GARANZIA, SIA QUELLA ESPLICITA, IMPLICITA O STATUTORIA, INCLUSO, MA NON LIMITATO A, QUALSIASI GARANZIA DI NON VIOLAZIONE DI PROPRIETÀ INTELLETTUALE, QUALITÀ SODDISFACENTE, IDONEITÀ A QUALSIASI SCOPO PARTICOLARE O QUALSIASI GARANZIA DERIVANTE DA PROPOSTA, SPECIFICAZIONI, CAMPIONI O ALTRO.

Questa garanzia non include la sostituzione di prodotti danneggiati a causa di abuso, incidente, uso inappropriato, negligenza, alterazione, riparazione, disastro, installazione o controllo inadeguati. Se il prodotto viene considerato difettoso per altri motivi, Intel, a sua discrezione, sostituirà o riparerà il prodotto, a proprie spese, eccetto nei casi qui sotto menzionati, a condizione che il prodotto venga consegnato congiuntamente al numero di autorizzazione per la restituzione del materiale (RMA, Return Material Authorization) (a) all'azienda da cui si è acquistato il prodotto, oppure (b) a Intel, solo quando in Nord America (se il prodotto è stato acquistato in Europa, sarà necessario consegnare il prodotto seguendo le modalità indicate in “(a)”). Se il prodotto viene inviato, il mittente si assume la responsabilità in caso di danni o di perdita durante il tragitto. È necessario utilizzare l'imballaggio originale del prodotto (o un suo equivalente) e pagare le spese di spedizione. Intel sostituirà o riparerà il prodotto (o la parte) con uno nuovo o uno rifabbricato, e il prodotto restituito diventerà proprietà di Intel. Intel garantisce che il prodotto riparato o sostituito sarà privo di difetti in materiale e in lavorazione per un periodo comunque non superiore: (i) a novanta (90) giorni dalla data di spedizione all'utente; oppure (ii) al periodo rimanente nella garanzia originale di tre (3) anni.

Questa garanzia dà all'utente diritti legali specifici; potrebbero esistere altri diritti, variabili da stato a stato. Tutte le parti e i componenti contenuti in questo prodotto sono coperti dalla garanzia limitata di Intel relativa a questo prodotto; il prodotto potrebbe contenere parti riciclate, completamente collaudate e garantite come nuove. Per maggiori informazioni sulla garanzia, chiamare uno dei numeri indicati qui sotto.

Restituzione di prodotti difettosi (RMA)

Prima di restituire un prodotto, contattare l'assistenza tecnica di Intel e richiedere un numero RMA; i numeri verdi sono qui sotto elencati:

Paese	Numero	Lingua
Francia	+33 (0) 1 41 91 85 29	Francese
Germania	+49 (0) 69 9509 6099	Tedesco
Italia	+39 (0) 2 696 33276	Italiano
Regno Unito	+44 (0) 870 607 2439	Inglese

Se il gruppo di supporto alla clientela determina che il prodotto è difettoso, richiederà l'emissione di un numero di autorizzazione per la restituzione del materiale (RMA) da porre all'esterno dell'imballaggio del prodotto. Intel non accetterà prodotti sprovvisti di tale numero visibile sull'imballaggio.

LIMITAZIONI DI RESPONSABILITÀ E RIMEDI

INTEL NON POTRÀ ESSERE CONSIDERATA RESPONSABILE DI ALCUN DANNO, DIRETTO O SPECULATIVO (INCLUSI, SENZA LIMITAZIONI COME INDICATO IN PRECEDENZA, I DANNI CONSEQUENZIALI, INCIDENTALI E SPECIALI) DERIVANTI DALL'USO O DALLA IMPOSSIBILITÀ DI UTILIZZARE QUESTO PRODOTTO, PER MOTIVI NON CONTEMPLATI NEL CONTRATTO, O DOVUTI A NEGLIGENZA, TORTO O SOTTO QUALSIASI GARANZIA, INDIPENDENTEMENTE DAL FATTO CHE INTEL SIA A CONOSCENZA O MENO DELLA POSSIBILITÀ DI TALI DANNI, INCLUSI, MA NON LIMITATI ALLA PERDITA D'USO, VIOLAZIONE DI PROPRIETÀ INTELLETTUALE, INTERRUZIONI D'AFFARI E PERDITA DI PROFITTI, NONOSTANTE QUANTO DETTO IN PRECEDENZA, LA RESPONSABILITÀ TOTALE DI INTEL NEI CONFRONTI DEI RECLAMI, SECONDO QUESTO ACCORDO, NON ECCEDERÀ IL PREZZO PAGATO PER IL PRODOTTO. QUESTE LIMITAZIONI SULLE RESPONSABILITÀ POTENZIALI SONO STATE FATTORE DECISIVO NELLA DETERMINAZIONE DEL PREZZO DEL PRODOTTO. INTEL NON ASSUME, NÉ AUTORIZZA ALCUNO AD ASSUMERE PER SÉ, NESSUN'ALTRA RESPONSABILITÀ.

Applicazioni di controllo di situazioni critiche: Intel disconosce specificatamente la responsabilità nel caso di uso dell'hardware in applicazioni di controllo di situazioni critiche (inclusi, al solo scopo di esempio, sistemi di controllo della sicurezza o della salute, dell'energia nucleare, o sistemi di controllo aereo o terrestre) da parte dei licenziatari o dei sottolicenziatari, e tale uso fa parte completamente del rischio intrapreso dall'utente. Il licenziatario è d'accordo nel difendere, indennizzare e liberare Intel da ogni reclamo risultante dall'uso del prodotto hardware in tale applicazioni da parte del licenziatario o del sottolicenziatario.

Software: il software accluso al prodotto hardware non è coperto dalla garanzia dell'hardware sopra descritta. Per maggiori dettagli sulla garanzia del software, vedere l'accordo di licenza relativo al software, inviato assieme al prodotto hardware.

Questa garanzia limitata dell'hardware è governata da, ed è conforme a, le leggi di Inghilterra e Galles. Il tribunale di Inghilterra avrà la completa giurisdizione su qualsiasi reclamo presentato sotto questa garanzia.

Beschränkte Hardwaregarantie (Nur für Europa)

Intel garantiert dem ursprünglichen Eigentümer, daß die in diesem Paket enthaltene Hardware keine Material- oder Herstellungsfehler aufweist. Diese Garantie gilt für drei (3) Jahre (a) nach dem Kaufdatum, wenn die ausgefüllte Registrierungskarte entsprechend den darauf enthaltenen Angaben zusammen mit einem Kaufnachweis eingesendet wurde; oder (b) nach dem Herstellungsdatum; oder (c) nach dem Registrierungsdatum, wenn die Registrierung innerhalb von 30 Tagen auf elektronischem Weg durchgeführt wird. Diese Garantie entfällt, wenn die Hardware bei der Installation beschädigt wird. Intel empfiehlt, die Installation durch den Verkäufer der Hardware durchführen zu lassen.

INTEL BEHÄLT SICH DAS RECHT VOR, IHREN AUFTRAG MIT EINEM PRODUKT ZU ERFÜLLEN, DAS NEUE ODER ERNEUERTE KOMPONENTEN ENTHÄLT. OBIGE GARANTIE GILT ANSTELLE ALLER ANDEREN AUSDRÜCKLICHEN, STILLSCHWEIGENDEN ODER GESETZLICH FESTGELEGTEN GARANTIEN. AUSGESCHLOSSEN SIND DAMIT AUCH UNTER ANDEREM ALLE GARANTIEN FÜR DIE VERKEHRSFÄHIGKEIT, DIE VERLETZUNG DER RECHTE VON DRITTEN, DIE EIGNUNG FÜR EINEN BESTIMMTEN ZWECK ODER GARANTIEN, DIE IM ZUSAMMENHANG MIT EINEM ANGEBOT, EINER SPEZIFIKATION ODER EINEM MUSTER GEGEBEN WURDEN.

Diese Garantie schließt den Hardware-Ersatz bei Beschädigung aufgrund von Mutwilligkeit, Unfall, falscher Verwendung, Fahrlässigkeit, Umänderung, Reparatur, Katastrophen, falscher Installation oder unvorschriftsmäßigem Testen aus. Wenn das Hardwareprodukt aus anderen Gründen beschädigt ist, liegt die Entscheidung bei Intel, ob die Hardware mit Ausnahme der im folgenden beschriebenen Einschränkungen kostenlos ersetzt oder repariert wird. Hierzu müssen Sie das Produkt zusammen mit einer Rückgabenummer (RMA-Nummer, siehe unten) entweder (a) an den Verkäufer des Produkts oder (b) an Intel zurücksenden (bei Kauf in Europa muß das Produkt an “(a)” geliefert werden). Das Risiko des Verlusts oder der Beschädigung während des Transports liegt bei Ihnen als Käufer. Sie müssen zum Versenden die Originalverpackung (oder einen gleichwertigen Ersatz) verwenden und die Versandkosten übernehmen. Intel ersetzt die Hardware entweder durch ein neues oder ein neuwertiges Produkt. Das zurückgegebene Hardwareprodukt wird Eigentum von Intel. Intel garantiert, daß das reparierte oder ersetzte Hardwareprodukt für einen Zeitraum von: (i) neunzig (90) Tagen ab Rückgabedatum oder (ii) für die verbleibende Zeit der ursprünglichen Garantie von drei (3) Jahren frei von Material- und Herstellungsfehlern ist. Dabei gilt jeweils der längere Zeitraum.

Mit dieser Garantie erhalten Sie bestimmte Rechte, die je nach Staat durch weitere Rechte ergänzt werden können. Alle Teile oder Komponenten dieses Hardwareprodukts werden durch die beschränkte Hardwaregarantie von Intel abgedeckt. Das Hardwareprodukt kann vollständig getestete, wiederverwendete Teile enthalten, die derselben Garantie wie neue Produkte unterliegen. Informationen zur Garantie erhalten Sie unter einer der Intel Kundendienstnummern, die am Ende dieses Handbuchs zu finden sind.

Rückgabe eines beschädigten Produkts (RMA)

Bevor Sie ein Hardwareprodukt zurücksenden, sollten Sie sich vom Intel Kundendienst eine sogenannte RMA-Nummer zuweisen lassen, indem Sie eine der folgenden gebührenpflichtigen Telefonnummern anrufen:

Land	Telefon	Sprache
Frankreich	+33 (0) 1 41 91 85 29	Französisch
Deutschland	+49 (0) 69 9509 6099	Deutsch

Land	Telefon	Sprache
Italien	+39 (0) 2 696 33276	Italienisch
Great Britain	+44 (0) 870 607 2439	Englisch

Nachdem die Beschädigung vom Kundendienst bestätigt worden ist, wird von der zuständigen Abteilung eine Rückgabenummer (RMA-Nummer) ausgegeben, die auf der äußeren Verpackung der Hardware angebracht werden muß. Intel akzeptiert kein Produkt ohne RMA-Nummer auf der Verpackung.

Haftungsbeschränkung und Rechtsmittel

INTEL HAFTET NICHT FÜR INDIREKTE ODER SPEKULATIVE SCHÄDEN (EINSCHLIESSLICH ALLER FOLGESCHÄDEN SOWIE ALLER ZUFÄLLIGEN UND BESONDEREN SCHÄDEN), DIE DURCH DIE VERWENDUNG ODER NICHTVERWENDBARKEIT DIESES PRODUKTS ENTSTEHEN, SEI DIES IM ZUSAMMENHANG MIT EINER VERTRAGLICHEN VERPFLICHTUNG, AUFGRUND VON FAHRLÄSSIGKEIT, DURCH UNERLAUBTE HANDLUNGEN ODER IM RAHMEN EINER GARANTIE. DIES GILT AUCH FÜR FÄLLE, IN DENEN INTEL ÜBER DIE MÖGLICHKEIT SOLCHER SCHÄDEN, DIE SICH UNTER ANDEREM DURCH NUTZUNGSAUSFÄLLE, BETRIEBSUNTERBRECHUNGEN UND GEWINNVERLUSTE ERGEBEN KÖNNEN, IN KENNTNIS GESETZT WURDE.

UNGEACHTET DER GEWÄHRTEN GARANTIE ÜBERSTEIGT DIE HAFTUNG VON INTEL IM RAHMEN DIESER VEREINBARUNG IN KEINEM FALL DEN KAUFPREIS DES HARDWAREPRODUKTS. DIESE HAFTUNGSBESCHRÄNKUNG IST EIN WESENTLICHER FAKTOR BEI DER FESTLEGUNG DES PREISES FÜR DAS HARDWAREPRODUKT. INTEL ÜBERNIMMT KEINE WEITERE HAFTUNG UND ERTEILT DRITTEN KEINERLEI BEFUGNIS, FÜR INTEL EINE WEITERE HAFTUNG ZU ÜBERNEHMEN.

Steuer- und Überwachungsanwendung von hoher Wichtigkeit: Intel schließt insbesondere die Haftung bei der Verwendung des Hardwareprodukts mit Steueranwendungen von hoher Wichtigkeit (z.B. Sicherheits- und Krankenversicherungssysteme, Steuersysteme für Nuklearanlagen sowie Verkehrsüberwachungssysteme für Boden- und Luftverkehr) durch den Lizenznehmer oder Unterlizenznehmer ab, und eine derartige Verwendung liegt ausschließlich in der Verantwortung des Benutzers. Der Lizenznehmer erklärt sich bereit, Intel zu verteidigen und schadlos zu halten bezüglich aller Klagen, die aus der Verwendung eines Hardwareprodukts für solche Zwecke vom Lizenznehmer oder Unterlizenznehmern erhoben werden.

Software: Die mit diesem Hardwareprodukt gelieferte Software wird von der oben beschriebenen Hardwaregarantie nicht abgedeckt. Bitte lesen Sie die entsprechende Softwarelizenzvereinbarung, die mit dem Hardwareprodukt geliefert wurde, um genaue Informationen zur Softwaregarantie zu erhalten.

Diese eingeschränkte Hardwaregarantie unterliegt den Gesetzen von England und Wales. Die englischen Gerichte sind Gerichtsstand für alle Klagen, die im Rahmen der Garantie erhoben werden.

Garantía limitada de hardware (sólo para Europa)

Intel garantiza al propietario original que el producto de hardware entregado en este paquete no tendrá defectos de materiales ni fabricación durante tres (3) años contados a partir de la fecha que resulte más reciente de entre las opciones siguientes: (i) la fecha de compra, sólo si devuelve la tarjeta de registro con prueba de compra de la forma indicada al respecto para registrarse; o bien (ii) la fecha de fabricación; o (iii) la fecha de registro, si éste se ha producido por medios electrónicos y dentro de los treinta (30) días siguientes a la compra. Esta garantía no cubre los daños sufridos por el producto durante el proceso de instalación. Intel recomienda que sea la empresa a la que adquirió el producto la que se encargue de su instalación.

INTEL SE RESERVA EL DERECHO DE CUMPLIMENTAR EL PEDIDO CON UN PRODUCTO QUE CONTENGA COMPONENTES NUEVOS O REFRABRICADOS. LA GARANTÍA ANTERIOR PREVALECE SOBRE CUALQUIER OTRA GARANTÍA, YA SEA EXPLÍCITA, IMPLÍCITA O REGLAMENTARIA, INCLUIDAS, SIN LIMITACIÓN, CUALESQUIERA GARANTÍAS DE NO INFRINGIMIENTO DE LA PROPIEDAD INTELECTUAL, CALIDAD SATISFACTORIA, ADECUACIÓN PARA UNA FINALIDAD DETERMINADA O CUALQUIER GARANTÍA SURGIDA DE CUALQUIER PROPUESTA, ESPECIFICACIÓN, MUESTRA O DE OTRA CLASE.

Esta garantía no cubre la sustitución de productos dañados por abuso, accidente, mal uso, negligencia, alteración, reparación, desastre, instalación incorrecta o pruebas incorrectas. Si el producto tuviera cualquier otro defecto, Intel se reserva la opción de reemplazar o reparar el producto sin cargo alguno, excepto los descritos a continuación, siempre que el producto se entregue con un número de autorización de devolución de material (RMA), a (a) la empresa a la que se adquirió o (b) a Intel, sólo en América del Norte (si lo adquirió en Europa, debe entregar el producto a "(a)". Si envía el producto, debe asumir el riesgo de daños o pérdida en el transporte. Debe utilizar el embalaje original (o equivalente) y costear los gastos de envío. Intel puede reemplazar o reparar el producto con piezas o productos nuevos o refabricados, y el producto devuelto pasa a ser propiedad de Intel. Intel garantiza que el producto reparado o reemplazado no tendrá defectos materiales ni de fabricación durante el periodo que resulte mayor de los siguientes: (i) noventa (90) días desde la fecha de envío; o (ii) el periodo de tiempo restante de la garantía original de tres (3) años.

Esta garantía le otorga derechos legales concretos y puede tener otros derechos que varían según la jurisdicción. Todas las piezas o componentes que contiene este producto están cubiertos por la garantía limitada de Intel sobre este producto; el producto puede contener piezas recicladas, completamente comprobadas, garantizadas como si de piezas nuevas se tratase. Si desea obtener más información sobre la garantía, puede llamar a uno de los números indicados a continuación.

Devolución de productos defectuosos (RMA)

Antes de devolver cualquier producto, póngase en contacto con el grupo de Asistencia al cliente de Intel y obtenga un número RMA en uno de los siguientes números no gratuitos:

País	**Número**	**Idioma**
Francia	+33 (0) 1 41 91 85 29	Francés
Alemania	+49 (0) 69 9509 6099	Alemán
Italia	+39 (0) 2 696 33276	Italiano
Reino Unido	+44 (0) 870 607 2439	Inglés

Si el grupo de Asistencia al cliente comprueba que el producto es defectuoso, se podrá en contacto con el Departamento de autorización de devolución de material para que éste le envíe un número RMA que debe colocar en el envoltorio externo del producto. Intel no puede aceptar productos sin el número RMA en el paquete.

LIMITACIÓN DE RESPONSABILIDAD Y REPARACIONES

INTEL NO SERÁ RESPONSABLE DE NINGÚN DAÑO INDIRECTO O ESPECULATIVO (INCLUIDOS, SIN LIMITAR A LOS ANTERIORES, LOS DAÑOS INDIRECTOS, INCIDENTALES Y ESPECIALES) PRODUCIDO POR EL USO O POR

LA IMPOSIBILIDAD DEL USO DE ESTE PRODUCTO, YA PROVENGA DE CONTRATO, NEGLIGENCIA, AGRAVIO O BAJO CUALQUIER GARANTÍA, SIN IMPORTAR QUE INTEL HAYA RECIBIDO PREVIO AVISO DE LA POSIBILIDAD DE DICHOS DAÑOS, INCLUIDOS, AUNQUE NO LIMITADOS A, PÉRDIDAS DE USO, INFRINGIMIENTO DE LA PROPIEDAD INTELECTUAL, SUSPENSIÓN DEL EJERCICIO COMERCIAL Y PÉRDIDA DE BENEFICIOS, A PESAR DE LO ANTERIOR, TODA LA RESPONSABILIDAD DE INTEL SOBRE LAS RECLAMACIONES REALIZADAS BAJO ESTE ACUERDO NO EXCEDERÁ EL PRECIO PAGADO POR EL PRODUCTO. ESTAS LIMITACIONES SOBRE LAS RESPONSABILIDADES POTENCIALES HAN CONSTITUIDO UN ELEMENTO ESENCIAL A LA HORA DE DETERMINAR EL PRECIO DEL PRODUCTO. INTEL NO ASUME NI AUTORIZA QUE NINGUNA PERSONA ASUMA EN SU LUGAR NINGUNA OTRA RESPONSABILIDAD.

Aplicaciones de control crítico: Intel deniega específicamente la responsabilidad por el uso del producto de hardware en aplicaciones de control crítico (incluidos, sólo a modo de ejemplo, los sistemas de seguridad o atención sanitaria, sistemas de control de energía nuclear o sistemas de control de tráfico aéreo o rodado) por Receptores o Subreceptores de la Licencia, y dicho uso queda enteramente a riesgo del usuario. El Receptor de la Licencia acuerda defender, indemnizar y mantener la inocencia de Intel por y contra toda reclamación surgida del uso del producto de hardware en tales aplicaciones por parte del Receptor o Subreceptor de la Licencia.

Software: El software proporcionado con el producto de hardware no está cubierto por la garantía de hardware descrita anteriormente. Si desea obtener información detallada sobre las garantías de software, consulte el acuerdo de licencia correspondiente al software incluido con el producto de hardware.

Esta garantía limitada de hardware se regirá e interpretará de acuerdo con las leyes de Inglaterra y Gales. Los tribunales de Inglaterra tendrán la exclusiva jurisdicción sobre todas las reclamaciones presentadas bajo esta garantía.

D Intel Customer Support

Intel offers a range of support services for your Intel® NetStructure™ 480T routing switch. You can learn about the options available for your area by visiting the Intel support Web site at http://www.intel.com/network/services.

Worldwide Access to Technical Support

Intel has technical support centers worldwide. The technicians who staff the centers generally offer service in the languages of the region.

Visit our Web site at http:/support.intel.com/.

North America only

For support, call **(800) 838-7136** or **(916) 377-7000**.

Japan only

For support, call **+81-298-47-0800**.

Other areas

For support in other countries, use the following table to dial the toll-free support number. Using the table, locate the country from which you are calling, dial the access number, await the dial tone, and then dial the listed 800 number.

Country	Dialing Information
Australia	1-800-881-011 await dial tone, then 800-838-7136
Austria [1][4]	022-903-011 await dial tone, then 800-838-7136
Belgium [1]	0-800-100-10 await dial tone, then 800-838-7136
China [3]	10811 await dial tone, then 800-838-7136
Denmark	8001-0010 await dial tone, then 800-838-7136
Finland [1]	9800-100-10 await dial tone, then 800-838-7136
France (includes Andorra)	19-0011 await dial tone, then 800-838-7136
Germany	0130-0010 await dial tone, then 800-838-7136
Hong Kong	800-1111 await dial tone, then 800-838-7136
India [5]	000-117 await dial tone, then 800-838-7136
Indonesia [2]	001-801-10 await dial tone, then 800-838-7136
Italy (includes Vatican City) [1]	172-1011 await dial tone, then 800-838-7136
Korea [1]	0-911 await dial tone, then 800-838-7136
Malaysia [4]	800-0011 await dial tone, then 800-838-7136
Netherlands [1]	06-022-9111 await dial tone, then 800-838-7136
New Zealand	000-911 await dial tone, then 800-838-7136
Norway	800-190-11 await dial tone, then 800-838-7136
Pakistan	0080001001 await dial tone, then 800-838-7136
Philippines	105-11 await dial tone, then 800-838-7136
Poland [1][3]	0-0-800-111-1111 await dial tone, then 800-838-7136
Portugal [3]	05017-1-288 await dial tone, then 800-838-7136
RSA (South Africa)	0-800-99-0123 await dial tone, then 800-838-7136
Russia [1][2][3]	755-5042 await dial tone, then 800-838-7136
Singapore	800-0111-111 await dial tone, then 800-838-7136
Spain	900-99-00-11 await dial tone, then 800-838-7136
Sri Lanka	430-430 await dial tone, then 800-838-7136
Sweden	020-795-611 await dial tone, then 800-838-7136
Switzerland [1]	0-800-550011 await dial tone, then 800-838-7136
Taiwan [1]	0800-10288-0 await dial tone, then 800-838-7136
Thailand [5]	0019-991-1111 await dial tone, then 800-838-7136
United Kingdom (BT) [3]	0800-89-0011 await dial tone, then 800-838-7136

Country	Dialing Information
United Kingdom (Mercury) [3]	0500-89-0011 await dial tone, then 800-838-7136
Vietnam	12010288 await dial tone, then 800-838-7136

Notes:

1 Public phones require coin deposit

2 Use phones allowing international access

3 May not be available from every phone

4 Public phones require local phone payment through the call duration

5 Not available from public phones

Index

A14542-001

100044-00 rev04